TP 10065E

A COMPUTER-ASSISTED X-RAY SCREENING SYSTEM

Prepared For

Transportation Development Centre
Policy and Coordination Group
Transport Canada

By

T. Feather, A. Lee-Kwen, L. Guan & I. Buchan
Array Systems Computing Inc.

and

M. Macecek & D. Allan
Techno-Scientific Inc.

JULY 1990

ARRAY SYSTEMS COMPUTING INC.
5000 Dufferin Street, Downsview, Ontario, M3H 5T5, Canada. Telephone: (416) 736-0900. Fax: (416) 736-4715. CNCP Mailbox: 21:ARY001

RECEIVED
SEP 4 1990
Transport Canada / Transports Canada

DOT 446609

CA1
DTC800-
90
C58 446610

91-03940/

41.D

Disclaimer

The contents of this report reflect the views of the authors and not necessarily the official view or opinions of the Transportation Development Centre of Transport Canada.

Un sommaire en francais de ce rapport est inclus avant la table des matieres.

Transport Canada / Transports Canada

PUBLICATION DATA FORM

1. Transport Canada Publication No.	2. Project No.	3. Recipient's Catalogue No.
TP 10065E	6918	

4. Title and Subtitle

A computer-assisted x-ray screening system

Final Report

5. Publication Date

July, 1990

6. Performing Organization Document No.

7. Author(s)

T. Feather, A. Lee-Kwen, L. Guan, I. Buchan, M. Macecek and D. Allan

8. Transport Canada File No.

D1455-168-3

9. Performing Organization Name and Address

Array Systems Computing, Inc.
5000 Dufferin Street, Suite 200
Downsview, Ontario
M3H 5T5

10. DSS File No.

11SD.T8200-7-7514-A

11. DSS or Transport Canada Contract No.

T8200-7-7514-01SD

12. Sponsoring Agency Name and Address

Transport Development Centre
200 Rene Levesque Blvd West
West Tower, Suite 601
Montreal, Quebec, H2Z 1X4

13. Type of Publication and Period Covered

Report

14. Project Officer

M. Audette

15. Supplementary Notes (Funding programs, titles of related publications, etc.)

16. Abstract

The objective of this project has been to develop a prototype real-time Computer-Assisted X-Ray Screening System (CAXSS) and demonstrate the feasibility of its use to assist operators of airport x-ray machines. CAXSS offers new computerized inspection tools and detection and highlighting of suspicious objects at a consistent performance level. The system interfaces easily with existing airport x-ray units. Guns and blades are the primary threats detected.

Detected threats in the displayed x-ray image are outlined in colour and announced with a beep. The operator may use touch-screen icons to examine suspicious areas by means of various image processing techniques. Transputers do most of the processing; a Personal Computer (PC) controls the system.

CAXSS meets the throughput goal of 12 bags per minute and tags about 15% of bags for hand search. The probability of detection (PD) for guns is about 90%; the false alarm rate (FAR) is 10%. The PD for blades is good; the FAR could be improved. Preliminary work on bomb components such as electronics and batteries shows promise.

We recommend continuation of threat detection efforts, field trials with an operational CAXSS, and integration of CAXSS with a dual-energy x-ray machine capable of discriminating organic material from inorganic for bomb detection.

17. Key Words	18. Distribution Statement
airport security, computer, image processing, threat detection, x-ray	Limited number of copies available from the Transportation Development Centre

19. Security Classification (of this publication)	20. Security Classification (of this page)	21. Declassification (date)	22. No. of Pages	23. Price
unclassified	unclassified		91	

TDC/CDT 79-005
Rev 88

446610

Canada

Transports Canada Transport Canada

FORMULE DE DONNÉES POUR PUBLICATION

1. N° de la publication de Transports Canada	2. N° de l'étude	3. N° de catalogue du destinataire
TP 10065E	6918	

4. Titre et sous-titre	5. Date de la publication
Un Systeme Logiciel de Détection aux Rayons X Rapport Final	Juillet, 1990
	6. N° du document de l'organisme exécutant

7. Auteur(s)	8. N° de dossier — Transports Canada
T. Feather, A. Lee-Kwen, L. Guan, I. Buchan, M. Macecek and D. Allan	D1455-168-3

9. Nom et adresse de l'organisme exécutant	10. N° de dossier — ASC
Array Systems Computing, Inc. 5000 Dufferin Street, Suite 200 Downsview, Ontario M3H 5T5	11SD.T8200-7-7514-A
	11. N° de contrat — ASC ou Transports Canada T8200-7-7514-01SD

12. Nom et adresse de l'organisme parrain	13. Genre de publication et période visée
Transport Development Centre 200 Rene Levesque Blvd West West Tower, Suite 601 Montreal, Quebec, H2Z 1X4	Rapport
	14. Agent de projet M. Audette

15. Remarques additionnelles (Programmes de financement, titres de publications connexes, etc.)

16. Résumé

Ce rapport concerne la mise au point d'un prototype de système de détection aux rayons X assisté par ordinateur fonctionnant en temps réel (Computer-Assisted X-Ray Screening System ou CAXSS), ainsi que les expérimentations visant à établir son utilité pour faciliter le travail des préposés aux appareils de détection par rayons X utilisés dans les aéroports. Le système CAXSS est un nouvel appareil électronique relié par interfaçage aux appareils à rayons X utilisés dans les aéroports et permettant de contrôler et de détecter les objets dangereux - notamment les armes à feu et les couteaux - avec un degré de fiabilité supérieur aux moyens actuellement disponibles sur le marché.

L'image des objets dangereux détectés par l'appareil à rayons X est affichée en couleurs et annoncée par un avertisseur sonore. A l'aide des icônes de l'écran tactile et de diverses techniques infographiques prévues à cette fin, le préposé peut examiner de plus près les zones où apparaissent des objets suspects. Le traitement parallèle est exécuté au moyen de transputers(ou transordinateurs) et le système est commandé par un micro-ordinateur.

Le système CAXSS contrôle les bagages à une cadence de 12 à la minute, et en désigne environ 15 p. 100 à la fouille systématique. Pour les armes à feu, le taux de probabilité de détection (PD) est d'environ 90 p. 100, tandis que le taux de fausse alarme (TFA) est de 10 p. 100. La probabilité de détection des couteaux est bonne mais le taux de fausse alarme pourrait être amélioré. Les résultats préliminaires concernant la détection des piles et des composants électroniques entrant dans la fabrication des bombes sont encourageants.

Il est recommandé de poursuivre le programme d'étude des moyens de détection des objets dangereux, de procéder à l'essai en aéroport d'un CAXSS opérationnel et de raffiner les techniques de détection des bombes en associant un système CAXSS à un appareil à deux sources radiogènes capable de distinguer les matières organiques des matières inorganiques.

17. Mots clés	18. Diffusion
la securité des aeroports, techniques infographiques, détection des objets dangereux, rayons X, ordinateur	Le Centre de développement des transports dispose d'un nombre limité d'exemplaires.

19. Classification de sécurité (de cette publication)	20. Classification de sécurité (de cette page)	21. Déclassification (date)	22. Nombre de pages	23. Prix
non classifié	non classifié		91	

Canada

Acknowledgements

We would like to acknowledge the support and help of the RCMP, the operations staff of Transport Canada at Pearson International Airport in Toronto, the Transportation Development Centre in Montreal, and Transport Canada personnel in Ottawa.

Executive Summary

1. PURPOSE OF THE PROJECT

Transport Canada (TC) is responsible for ensuring that potentially dangerous items such as guns and bombs are not carried on board airplanes departing from Canadian airports. Various screening techniques are currently being used by TC, and others are being developed. It is customary to x-ray all carry-on baggage for this purpose. Any bag suspected of containing a "threat" is subjected to a hand search.

The interpretation of x-ray images is often difficult, and requires a high level of concentration that is hard to maintain.

The objective for this project is to develop a prototype real-time Computer-Assisted X-ray Screening System (CAXSS). The system is intended to assist operators of airport x-ray screening machines by offering new inspection tools such as zooming and image enhancement, and by detecting and highlighting suspicious objects at some reasonable consistent performance level. It should perform better than an inattentive operator. However, present technology is unlikely to match the threat recognition ability of an alert, fully-trained operator, particularly at real-time speeds and reasonable cost. The system is therefore not intended to replace human operators.

2. PROJECT SCOPE

A prototype real-time system has been built. It was designed to interface with an operational x-ray screening unit as easily as possible (with minimal changes to existing airport hardware). Two major milestones were defined: a "basic" system, and an "advanced" system. Testing on the basic system provided feedback for the advanced system.

The purpose of the prototype is to prove clearly the feasibility of a real-time system with adequate detection performance at an acceptable cost. Goals are a 90% detection rate and a 10% false alarm rate, with average throughput of 12 bags per minute.

The two threat categories addressed are guns, exposed or hidden, and blades such as knives and scissors; Some preliminary work has been done on the detection of bomb components - electronics and batteries.

3. SYSTEM DESIGN

3.1 Hardware

There are two categories of hardware: the interface to the existing SCANRAY Linescan System One x-ray machines, and the processing hardware.

The basic processing hardware consists of:

- an 80386 PC,
- a 19-inch colour monitor with touch screen,
- an ATI Super-VGA Wonder Video Card,
- an Inmos T-800 transputer board with up to four transputers.

This hardware configuration was chosen on the basis of cost, real-time processing constraints, ease of expansion and probable portability of the developed software. Most of the software for both PC and transputers is written in C.

The interface to the x-ray machine consists of:

- a custom interface board in the Linescan unit,
- a custom digital buffer board in the PC.

3.2 User Interface

An operator deals with CAXSS via a single screen. The current x-ray image is displayed on this screen. Threats are highlighted as they are detected automatically, and an alarm beep may sound. The operator may use the the touch-screen icons on the right-hand side of the display to inspect more closely areas that seem suspicious.

3.3 Software

The PC is mainly responsible for acquiring data from the x-ray interface, sending the data to the transputers, receiving threat descriptions back from the transputers, and highlighting threats on the screen. The four transputers handle most of the image processing and threat analysis. One is dedicated to blade detection. The others search for guns and other dense threats.

Preliminary work on electronics detection involves texture analysis.

4. RESULTS

To demonstrate the system's detection ability, a set of test images was constructed to measure the probability of detection (PD). Airport bags were examined to measure the false alarm rate (FAR).

CAXSS meets the throughput goal of 12 bags per minute and tags about 15% of bags for hand search. The probability of detection (PD) for guns is about 90%; the false alarm rate (FAR) is 10%. The PD for blades is good; the FAR could be improved. Preliminary work on bomb components such as electronics and batteries shows promise.

The operator interface proved successful, with only a few minutes training required for new operators to learn how to use the touch screen. The coloured highlighting of threats was found to be effective, and the interactive processing features useful for examining suspicious objects.

5. CONCLUSIONS

The prototype CAXSS demonstrates the feasibility of supporting x-ray machine operators in airports with an automatic threat analysis system.

Overall

1. Average throughput of CAXSS meets the goal of 12 bags per minute.
2. The search rate of about 15% matches that of human operators.
3. The operator interface is simple and useful.
4. The prototype can easily be converted to operational use.

Gun Detection

1. The PD for guns is good (91% at the yellow and red threat levels, vs 85% for human operators).
2. The FAR is acceptable; almost all false alarms were for suspicious-looking objects that should be hand searched.

3. The segmentation could be improved.

Blade Detection

1. The PD for blades is good.

2. The FAR could be improved.

6. RECOMMENDATIONS

1. We recommend the development of a complete operational system to be used for field trials.

2. We recommend that development should continue, to improve the current system, and to detect bomb components such as electronics, batteries, wires and possibly detonators.

3. We recommend that the system be integrated with a dual-energy x-ray machine capable of discriminating organic material from non-organic. This would provide a critical additional component in a comprehensive bomb recognition model.

SOMMAIRE

1. BUT

Transports Canada (TC) a pour mission de s'assurer qu'aucun objet dangereux, telles des bombes ou des armes à feu, ne soit transporté à bord des avions, aux aéroports canadiens. Pour cela, le Ministère fait appel à divers moyens de détection - dont l'examen par rayons X - pour identifier les bagages devant faire l'objet d'une fouille systématique. Il s'intéresse donc vivement aux programmes de recherche en matière de contrôle des bagages transitant dans les aéroports.

Or, l'interprétation des images produites par les rayons X est souvent difficile et requiert de la part des préposés un degré de concentration difficile à maintenir sur une période de temps prolongée.

Le but de ce projet est de mettre au point un prototype de détecteur aux rayons X fonctionnant en temps réel et assisté par ordinateur (CAXSS). Grâce à cet appareil, les préposés à la détection des objets dangereux pourront grossir et mettre en évidence les images obtenues par rayons X, et seront mieux à même de détecter les objets suspects. Précisons toutefois que si le CAXSS est plus performant qu'un préposé distrait, dans l'état actuel de la technique, il ne saurait égaler le rendement, en termes de rapidité et de coût-efficacité, d'un préposé vigilant rompu aux techniques de détection. Ce système ne se définit donc aucunement comme un moyen de remplacer les opérateurs humains.

2. DESCRIPTION

Un prototype fonctionnant en temps réel a été construit. Cet appareil a été conçu pour s'adapter aisément à un détecteur par rayons X (c'est-à-dire avec le minimum de modifications au matériel actuellement mis en oeuvre dans les aéroports). Le projet consistait, dans un premier temps, à mettre au point un système de base et, dans un deuxième temps, à mettre au point un système "évolué". Les essais effectués sur le système de base ont servi de point de départ à la conception du système évolué.

L'objet du prototype était d'établir qu'il est possible de réaliser un système fonctionnant en temps réel tout en obtenant un taux de détection adéquat à un coût raisonnable. Les objectifs visés étaient un taux de détection de 90 p. 100 et un taux de fausse alarme de 10 p. 100, compte tenu d'une cadence de 12 bagages à la minute.

Le système étudié sert principalement à la détection des armes à feu et des couteaux. Une étude préliminaire a également été réalisée concernant la détection des piles et des composants électroniques entrant dans la fabrication des bombes.

3. CONCEPTION DU SYSTEME

3.1 Matériel

Le système est constitué de deux catégories de matériel: le module d'interface avec les appareils de détection par rayons X, SCANRAY Linescan System One, et le module infographique. Le module d'interface, de type numérique, a été étudié exprès.

Le matériel de traitement de base est constitué des éléments suivants:

- un PC 80386,
- un moniteur en couleurs, de 48 cm à écran tactile,
- un module vidéo ATI Super-VGA Wonder
- un transputer (ou transordinateur) Inmos T-800 pouvant contenir jusqu'à 4 puces.

Les critères qui ont présidé au choix de cette architecture sont le coût, les contraintes liées au traitement en temps réel, les possibilités d'extension du système et la portabilité du logiciel. Dans l'ensemble, le logiciel d'exploitation du PC et des transordinateurs a été rédigé en langage C.

3.2 Interface utilisateur

L'opérateur dialogue avec le système CAXSS par l'intermédiaire d'un seul écran affichant les images produites par rayons X. Dès leur détection par le système, les objets suspects sont mis en évidence à l'écran et un signal sonore se fait entendre. L'opérateur peut alors examiner de plus près la zone suspecte au moyen des icônes tactiles situées à la droite de l'écran ou des fonctions infographiques prévues à cette fin. Au besoin, il peut faire appel au menu d'aide en ligne.

Les objets suspects sont analysés et on leur attribue l'une des trois cotes de danger désignées par un code de couleur distinct.

3.3 Logiciel

Les principales fonctions du PC sont de recueillir les données provenant du module d'interface avec le détecteur aux rayons X, de transmettre ces données aux transputers Inmos, de recevoir les descriptions d'objets dangereux produites par ces derniers et de mettre en évidence les images correspondantes. Ce sont les quatre transputers - dont l'un est affecté à la détection des couteaux, les trois autres aux armes à feu et autres objets dangereux de haute densité - qui exécutent la plus grande partie du traitement des images et de l'analyse des objets dangereux.

Les travaux préliminaires visant la détection des composants électroniques intéressent l'analyse des textures.

4. RESULTATS

Pour contrôler le rendement du système, on a mesuré la probabilité de détection (PD) à partir d'une série d'images témoins. Ensuite, des bagages ont été examinés pour évaluer le taux de fausse alarme (TFA).

Le système CAXSS répond à l'objectif de 12 bagages à la minute et désigne environ 15 p. 100 des bagages à la fouille systématique. Dans le cas des armes à feu, la probabilité de détection (PD) est d'environ 90 p. 100 tandis que le taux de fausse alarme (TFA) est de 10 p. 100. La PD des couteaux est satisfaisante; le TFA pourrait toutefois être amélioré. Les résultats préliminaires concernant la détection des piles et des composants électroniques entrant dans la fabrication des bombes sont encourageants.

L'interface utilisateur est convivial: il ne faut que quelques minutes de formation pour apprendre à un opérateur comment utiliser l'écran tactile. La méthode de mise en évidence en couleurs des objets dangereux s'est avérée efficace et les fonctions de traitement interactives facilitent l'examen des objets suspects.

5. CONCLUSIONS

Le prototype CAXSS a permis d'établir l'intérêt de doter les opérateurs d'appareils aux rayons X mis en oeuvre dans les aéroports d'un système d'analyse automatique des objets dangereux.

Sommaire

1. La cadence moyenne atteinte par le système CAXSS répond à l'objectif de 12 bagages/minute.

2. Le taux de renvoi à la fouille systématique correspond, à 15 p. 100 à celui atteint par les préposés.

3. L'interface utilisateur est convivial

4. Le prototype peut aisément être converti en un système opérationnel.

Détection des armes à feu

1. La PD des armes à feu est satisfaisante (91 p. 100 des cas relevant des cotes de danger jaune et rouge par rapport à 85 p. 100 pour les préposés).

2. Le TFA est acceptable, presque toutes les fausses alarmes sont attribuables à des objets suspects qui justifient une fouille systématique.

3. La segmentation de l'image pourrait être améliorée.

Détection des couteaux

1. La PD pour les couteaux est satisfaisante.

2. Le TFA pourrait être amélioré.

Détection d'autres objets dangereux

1. Il faudrait poursuivre les études concernant la détection des piles et des composants électroniques entrant dans la fabrication des bombes.

6. RECOMMANDATIONS

1. La mise au point d'un système entièrement opérationnel aux fins d'essai en aéroport;

2. La poursuite des travaux visant à perfectionner les systèmes existants et à améliorer la détection des composants électroniques, piles, fils et détonateurs entrant dans la fabrication des bombes;

3. L'intégration à un appareil à deux sources radiogènes capable de distinguer les matériaux organiques des matériaux inorganiques, pour faciliter la mise au point d'un système capable de détecter les bombes.

Table of Contents

1. INTRODUCTION 1

1.1 Background 1
1.2 Project Objectives 2

2. PROJECT PLAN 4

3. SYSTEM DEFINITION 5

3.1 Technology Review 5

3.1.1 Image Processing Environment 5
3.1.2 X-ray Machines 6
3.1.3 Literature Review 8

3.2 Threat Definition 8

3.2.1 Threat Categories 8
3.2.2 Threat Detectability 9

3.3 System Specification 11

3.3.1 Threat Identification 11
3.3.2 System Performance Measures 12
3.3.3 System Throughput 15
3.3.4 Basic System 15
3.3.5 Advanced System 15

4. SYSTEM DESIGN 17

4.1 Hardware Overview 17

4.1.1 Digital Interface to Linescan One 18

4.2 Software Overview 19

4.2.1 Introduction 19
4.2.2 Input from X-ray Machine 19
4.2.3 Processing on PC 20
4.2.4 Processing on Transputers 20

4.3 User Interface 21

4.3.1 Operator Interface 22
4.3.2 Supervisor Interface 24
4.3.3 Analyst Interface 26

5. THREAT RECOGNITION ALGORITHMS 29

5.1 Introduction 29
5.2 Handguns 29

5.2.1 Object Segmentation 29
5.2.2 Feature Extraction 32
5.2.3 Object Classification 35

5.3 Knives and Scissors 39

5.3.1 Image Pre-Processing 39
5.3.2 Object Segmentation 40
5.3.3 Feature Extraction 40
5.3.4 Object Classification 42

5.4 Other Threats 44

5.4.1 Electronics 44
5.4.2 Batteries 45

6. RESULTS 46

6.1 Threat Detection 46

6.1.1 Guns 46
6.1.2 Knives 49
6.1.3 Other Threats 51
6.1.4 Overall Alarm Rate 52

6.2 Processing Speed 52

6.2.1 Airport Bags 53
6.2.2 Test Bags Containing Guns 54
6.2.3 Bags Containing Knives and Scissors 55

6.3 User Interface 55
6.4 Developed Software 56

7. CONCLUSIONS 57

8. RECOMMENDATIONS 59

REFERENCES 60

GLOSSARY OF TERMS 61

BIBLIOGRAPHY 62

Appendix A: Project Schedule and Documents

Appendix B: CAXSS System Parameters

Appendix C: Image Processing Techniques

List of Figures

4.1 CAXSS Prototype System Configuration 17

5.1 Flat Non-occluded Guns 31
5.2 Flat and Horizontal Guns 33
5.3 Non-occluded Knives 41
5.4 Blades in Briefcase 43

1. INTRODUCTION

1.1 Background

Transport Canada (TC) is responsible for ensuring that potentially dangerous items are not carried on board airplanes departing from Canadian airports. Such items include guns, bombs, and other weapons. Various screening techniques are currently being used by TC, and others are being developed, to ensure that passengers do not carry such "threats" onto an airplane.

One of the primary screening techniques is the use of an x-ray machine to examine the contents of passengers' carry-on baggage. All carry-on baggage is put through an x-ray screening device and an operator inspects the resulting x-ray image. If the operator decides that a bag might contain a threat, the bag is subjected to a hand search.

Although x-ray screening can be very effective, it has several drawbacks. These can be divided into two main classes: limitations of the x-ray imaging process, and problems with image interpretation. The x-ray imaging process has many limitations, the major one being that x-ray imaging provides good definition only of relatively dense materials. While this is not a problem for a threat such as a handgun, various critical components of an explosive device may not be visible at all. In addition, object recognition is adversely affected by occlusion by other objects of similar or higher density. It can be very difficult to identify overlapping objects, and not-too dense objects can be easily hidden.

Even when it is possible to obtain a good x-ray image of a threat, interpretation of the image to identify that threat presents many difficulties. A major problem is that a high degree of concentration and thoroughness is required. Since objects are x-rayed with transmitted rather than reflected radiation, the image represents the mass of the object rather than its familiar surface. Self-occlusion, whereby important parts of an object may be masked by other denser parts of the same object, can pose a serious recognition problem. The reliable interpretation of x-ray imagery requires operators with training and experience. Even so, in a high-volume security screening station, it can be very difficult for an operator to maintain a sufficiently high level of concentration for every bag.

In the last few years, computer vision technology has progressed to the point where it may be cost-effective to apply it in an attempt to enhance this x-ray screening process. Recognizing this, the Transportation Development Centre (TDC) awarded a contract to Roy Ball Associates (RBA) to develop a prototype automated threat detection system with the goal of proving the feasibility of such a system. Some useful information emerged, but the overall results of this contract were inconclusive. The RBA system attempted to detect handguns and one particular bomb, as well as dark regions that could hide dense threats such as a handgun. Detection of dark regions, or "blobs", was, by definition, 100%. The detection rate on a very small sample of handgun images was 72%; the false alarm rate, also on a very small sample, was 17%.

The RBA system was also able, not surprisingly, to detect the only bomb that was studied. The method of detecting this bomb was so specific that it did not really provide any information about the feasibility of finding bombs in general.

The detection performance for handguns did, however, demonstrate the likely feasibility of acceptable detection and false alarm rates for handguns. However, the RBA system was extremely slow, being more than two orders of magnitude slower than a real system would have to be for handguns and blobs. Thus it was not clear that a real-time system could be developed with present generation hardware at an acceptable cost.

RBA also investigated the possibility of modifying the x-ray machines to incorporate either a dual-angle or dual-energy capability; however, no conclusive results were achieved.

On-going work at Carnegie-Mellon University and the University of North Carolina has yielded similarly inconclusive results.

In May 1988, following a competitive tender, a contract was awarded by TDC to Array Systems Computing (ASC) and Techno Scientific Inc (TSI) to develop a prototype real-time Computer-Assisted X-ray Screening System (CAXSS).

1.2 Project Objectives

The main objective for this project is to develop a prototype real-time system capable of assisting operators of airport x-ray screening machines by detecting and highlighting dangerous objects.

The prototype should prove clearly the feasibility of producing a real-time system with adequate detection performance at an acceptable cost. Goals are for a 90% detection and a 10% false alarm rate with a throughput rate of 12 bags per minute. The threats to be detected include handguns, and dark regions that could conceal, for example, handguns and possibly some types of bombs. The prototype system should work with TC's standard x-ray machine, which is the SCANRAY Linescan System One manufactured by EG&G Astrophysics.

A secondary objective is for the prototype to be developed in such a way that it could be re-engineered as easily as possible into an operational system.

It is important to note that it is not an objective of the system to attempt to replace human operators. This would be a completely impractical goal. Rather, the system is intended to enhance an operator's performance. It is clear that the system cannot, in general, recognize objects as well as a human being can. This is particularly true in cases of partial occlusion. Although the system may never be as good as a fully attentive, or "ideal", human operator, it should be able to provide a very consistent minimum level of performance. In this way, the system can improve on an operator's performance in inevitable instances of fatigue leading to lapses in attention and inconsistency.

2. PROJECT PLAN

The project personnel are structured into two teams: the image processing team (ASC) and the x-ray team (TSI). ASC is primarily responsible for the system design and development, including the image processing software and hardware. TSI is primarily responsible for interfacing the system to the x-ray machine and for examining possible changes to the x-ray hardware.

It was decided early in the project to concentrate on software first in order to make the best possible use of the information that was presently available. Changes to the x-ray hardware were to be considered only when required by the software. The motivation for this includes the existence of a large installed base of Linescan machines, which cannot be retrofitted easily without substantial time, cost, and disruption of operations. Ideally, CAXSS should require no changes at all to the x-ray machines and should interface to the Linescan in the simplest possible manner.

The development methodology chosen was rapid prototyping - to put together a relatively complete real-time system as quickly as possible, and then to refine it iteratively. Two major milestones were defined: a "basic" system, and an "advanced" system. The basic system is a complete, fully-functional prototype that could be delivered as soon as possible in order to allow TC to test the system and provide feedback for the advanced system.

A detailed work plan was produced as the first contract deliverable. The project schedule and list of deliverable documents are given in Appendix A. There are two particularly significant documents for users of the system. The System User Documentation manual gives details of the hardware configuration and provides complete instructions for the user interface. The Maintenance Manual describes the software in sufficient detail to allow a reader with programming experience to maintain and modify CAXSS.

3. SYSTEM DEFINITION

3.1 Technology Review

3.1.1 Image Processing Environment

An extensive market review of existing image processing hardware was performed. The critical evaluation criteria were the real-time processing constraints, the degree of expandability of the configuration, the cost, and the likely degree of portability of the developed software.

A PC chassis with embedded processing boards was chosen as the best architecture. The main reason for this was low cost and ease of expandability. It was determined that a standard AT bus could easily meet the required data transfer rates. An 80386 processor was also considered to be more than adequate to handle any required pre-processing of the data.

It was decided to use multiple general-purpose processors as the main computational engine. Specifically, the Inmos T-800 transputer was chosen. The transputer is specifically designed for parallel processing applications. Our review showed that the transputer offered the best price-performance ratio, software portability via a C compiler and tremendous ease of expansion. System capability can be expanded by adding extra transputers, and/or by upgrading as new generations become available.

We also reviewed and benchmarked the Occam language compiler for the transputer. Occam was designed specifically for parallel processing using transputers. However, we found the code generated by the Occam compiler to be exceptionally poor. It appeared to perform almost no optimizations, resulting in very high overhead for loops and subscripted arrays. Benchmarks showed that this overhead typically dominates the execution time. Furthermore, Occam does not include the concept of pointers so that there is no way to circumvent this considerable overhead. The use of Occam would also greatly reduce the portability of the software. For these reasons, Occam was rejected.

Implementing the software, both for the 80386 and the transputers, in C allows the software to be ported to other processors relatively easily. It also allows algorithm development and testing to be performed on any system with a suitable C compiler.

For the prototype, Inmos' modular transputer boards were chosen for their flexibility. One of these boards holds up to 4 transputers, each with 2 MB of memory. Inmos rates the performance of each 20 MHz transputer at 10 Million Instructions Per Second (MIPS), and the newer 25 MHz version (which became available during the project) at 12.5 MIPS. The T-800 also has a high speed floating point processor on the chip. Multiple boards can be easily connected via high-speed transputer links, bypassing the PC bus, which would only have to be used for host I/O.

It was assumed that the basic system would employ two transputers, and the advanced system four, based on cost and software considerations.

Many other possible processors were examined and rejected. In particular, various array processors, both floating-point and bit-serial, were rejected primarily for price-performance reasons, but also because of difficulty of expansion and lack of portability of the developed software.

Several display boards were considered. It was decided to use some standard memory-mapped board, rather than to commit to a particular specialist board.

3.1.2 X-ray Machines

The Linescan x-ray machine was studied in detail. The objective of the x-ray studies was to gain a good understanding of the physical phenomena involved, the x-ray hardware used, and the physical limitations of the system. A homogeneous, quantifiable x-ray field is necessary if pattern recognition software is to have any chance of success.

The studies included:

- x-ray beam geometries,
- object image distortions,
- inhomogeneities across the diode row,
- response of the machine and setting of the dynamic range,
- calibration of the unit.

A series of threats including handguns and explosive devices were x-rayed. The RS-170 output of the x-ray machine was digitized via a frame grabber and the resulting digital images saved for analysis.

The general quality of the images was poor. They exhibited substantial variations in pixel values over homogeneous materials, which were apparently random in nature. The system also seems to be very sensitive to fluctuations in the voltage of its power source. The machine periodically performs a self-calibration on each of its diodes, but this does not seem to be very effective with the System One input board. It can only calibrate within a limited dynamic range, with the result that many diodes are (randomly) not calibrated. This tends to happen more towards the edges of the belt, where objects are seen less frequently.

A partial explanation for the poor image quality can be found in an examination of the steps between acquisition and display. The x-ray system actually produces an 8-bit (maximum) digital image in its internal memory. This is then converted back to the analogue RS-170 signal, which then has to be re-digitized for storage, and finally converted back to analogue for display. These multiple conversions increase the noise in the signal, change the linearity and dynamic range, and degrade the picture quality.

It was decided that a digital interface to the x-ray should be built in order to avoid the unnecessary A/D conversions. A digital interface also greatly simplifies the processing of data since each line scanned can be transferred to the interface immediately, to be passed to the processor at once or in buffers of any convenient size. With an analogue interface, exactly 512 lines at a time would be available, displaced from the previous 512 by some variable amount, dependent on whether and for how long the belt had been stopped.

X-ray intensities for objects of different opacities were quantified using specially built calibration pieces of precisely ground steel - bars about 4 mm thick and a wedge from 0 to 5 mm. Any thickness of steel, up to the maximum 13 mm penetration specified for the Linescan System One machine, could be tested with combinations of these pieces.

In addition, a number of experiments examining the x-ray source itself were carried out, using x-ray film as the recording medium to avoid non-linearity of the diode system.

Images of collimated slits were exposed and scanned with a microdensitometer to produce an x-ray intensity profile across the belt. The scans revealed significant inhomogeneity in the intensities (see Figure 3.1).

While these peaks and valleys are partly compensated for by the correction circuit of the input board, the amplitude often exceeds the dynamic range of this correction, in which case no correction is made. The problem is compounded by the fact that these intensity variations are time-dependent.

Several possible modifications to the x-ray system were also considered and costed. These were:

- variable voltage system,
- E-Scan dual-voltage system,
- stereo or dual-head system,
- higher resolution detectors, and
- Z-Scan backscatter system.

Since the cost of each of these was significant, it was decided not to proceed with any modifications until the data from the existing single-head system had been fully exploited.

3.1.3 Literature Review

A review of applicable literature was undertaken. A large number of references was compiled; see the Bibliography. No literature on "similar" systems could be identified. Although a large number of real vision systems do exist, most work on very limited domains of a few known objects only.

3.2 Threat Definition

3.2.1 Threat Categories

A detailed threat definition document was prepared, cataloguing known threats that might be carried onto an aircraft. A large sample of such threats was provided by the RCMP and was x-rayed. The resulting images were photographed, and stored on videotape and in digital format. The primary threat categories are summarized below.

Table 3.1 Primary Threat Categories

1. Hand guns (revolvers, autoloaders)
2. Other firearms
3. Toy guns
4. Sharp instruments (knives, scissors, etc.)
5. Improvised explosive devices
   (explosive material + battery + wires + detonator)
6. Grenades
7. Pipe bombs
8. Bottles containing liquid
9. Aerosol cans
10. Blunt instruments (clubs, etc.)

Many other categorizations are possible. This one tries to group threats that have a similar appearance.

An additional "synthetic threat" category is that of Dark Connected Regions (DCRs) that could hide a dense threat, such as a gun. This subjective definition depends heavily on the ability of the observer (man or machine) to segment and interpret an image. Naturally, DCRs constructed by disguising real threats such as handguns should clearly be considered threats and flagged by the system. Such constructed DCRs could easily be used for objective testing of the detection rate, but any measurement of the false alarm rate, requiring an answer to the question "Could this object hide a threat?", would obviously be much more subjective, given the diversity of threat sizes and densities. It makes more sense to associate each DCR found with a particular real threat, and categorize it as an example of occlusion (up to 100 percent).

3.2.2 Threat Detectability

The threat images were analyzed to determine the feasibility of automatic detection.

Images taken of handguns are readily recognizable when the handguns are in a flat position. In other orientations, when the handle of the gun cannot be seen, the gun is much less recognizable; however, in cases where the barrel is visible (i.e., when it is roughly perpendicular to the x-ray beam), the gun is still

distinguishable from objects typically seen in handbags. When the gun barrel is pointed up towards the x-ray beam, it is much more difficult to recognize the gun. In general, a gun in this orientation cannot reliably be distinguished from typical dark items in handbags.

Toy guns made of light plastic are very difficult to recognize due to their lack of dense components. Even very mild occlusion can make them unrecognizable.

Some other dense threats such as occluded guns and pipe bombs, although not readily recognizable on the basis of shape are generally distinguishable from objects typically found in bags based on their size and density.

Knives and scissors, although less dense, are also generally recognizable when they are not occluded and are lying flat. Their lower density unfortunately makes them very susceptible to occlusion by objects that are typically found in carry-on bags. When the blade is parallel to the x-ray beam, or is moderately occluded, they are generally not recognizable. In some of the latter cases, when the knife is sufficiently large or dense, it may be distinguishable with a reasonable false alarm rate on the basis of its size and density.

Explosive devices are very difficult to categorize. However, in general they have at least four components: explosive material, a battery, wires, and a detonator. They are therefore best disguised by items that naturally have wires and batteries, such as a radio or other electronic items. A bomb may also have a timing device, which could also be electronic, such as a calculator or watch.

Small size and low density makes detonators are difficult to recognize, even when not obstructed by any other object. Inside a typical bag, they are well camouflaged. Plastic explosive material is visible if there is a sufficient thickness, but has no general feature that would differentiate it from a multitude of commonly occurring materials. Batteries, if they are standard 9-volt, C, D, AA, or AAA cells, can typically be seen, but with their small size and lack of distinguishing features, they are difficult to recognize and highly subject to occlusion and confusion with other small dense objects.

Wires, of the gauge required in a bomb, are visible in images, but wires, and wire-like streaks, are extremely common in typical bags. There may be some hope, however, of detecting wires attached to non-occluded batteries. Some electronic devices such as radios are

often recognizable in x-ray images from the complex texture that they exhibit. Electronic timing devices, such as a typical wrist watch or calculator, tend to be too small or too light to recognize.

Bottles filled with liquid, if sufficiently large, appear to be quite dark in x-ray images. Even though their shape can vary widely, their density, combined with size and shape constraints makes them potentially detectable.

Finally, aerosol cans may be detectable if they are not occluded by other objects. Generally, their relatively low density makes occlusion likely.

3.3 System Specification

3.3.1 Threat Identification

It was decided to employ three levels of threat highlighting/alarms defined very roughly as the following:

Table 3.2 Threat Levels

Level	Colour	Approximate Meaning	Recommendation
3	RED	Shows many characteristics of a threat	Search bag
2	YELLOW	Shows some characteristics of a threat	Inspect object and search bag unless clearly not suspicious
1	BLUE	Could conceal a threat	Inspect object and search bag if it appears suspicious
0	-	Probably not a threat	

A level 0 threat is not usually highlighted in any way (although it can be optionally highlighted in brown for debugging purposes).

In general, a red alarm is meant to indicate a definite

search. A yellow alarm indicates that a search should be conducted unless the operator is certain that the object is not dangerous. A blue alarm indicates that the operator should examine the area carefully and search the bag only if the area appears to be suspicious. The probability-of-detection and false-alarm-rate goals are therefore primarily relevant at the red and yellow levels.

An identified threat is to be highlighted with an outline in the colour corresponding to the threat level. Outlining a threat was chosen in favour of shading so that the object could be easily inspected without having to remove the highlighting in most cases. A beep may also sound.

3.3.2 System Performance Measures

The basic detection goal was set at a 90% Probability of Detection (PD) with a 10% False Alarm Rate (FAR). We also compute an "efficiency" measure, defined as

$$\text{Efficiency} = \frac{\text{Log}(\text{FAR})}{\text{Log}(\text{PD})}$$

which we consider to be a single reasonable measure of the overall performance of the system. The rationale for this is that PD and FAR tend to be exponential functions of each other, so that efficiency tends to be constant, at least over some limited range.

A PD of 90% and FAR of 10% correspond to an efficiency of 21.8, which is approximately the level we attempted to achieve. Efficiency is useful for comparing PD and FAR results between different classes and threat levels.

PD, FAR, and efficiency are measured for each threat level based on that level plus all higher levels. For example, the PD at the blue level is the percentage of all threats coloured blue, yellow, or red.

Initially the PD and FAR were defined relative to the performance of an ideal operator; however, this was impractical to measure with any objectivity, so absolute definitions of the PD and FAR rates were adopted, as defined below.

The PD was simply defined as the percentage of bags containing real threats that were detected by the system. A PD is defined for each well-defined threat class or sub-class. Ideally sub-classes are defined so that the system performs similarly on all members of that class.

This results in minimal dependency on the particular sample used to measure the PD. For example, it is of little use to compute a PD over guns and knives combined, since there is no meaningful way to define a "typical" mixture of images containing guns and knives. The principle extends from this obvious case to subclasses of each type of threat, which can be very different in terms of their detectability (e.g., guns in different orientations).

The FAR was defined to be the percentage of real bags passing through an airport x-ray machine for which the system detects a possible threat. (The percentage of bags containing a real threat is sufficiently small that it can be ignored.) This figure is difficult to interpret since the FAR, under this definition, should not be 0%; an object that could obscure a threat should cause an alarm. Objects that could conceal a handgun, for example, regularly appear in typical carry-on baggage. A relative FAR measure that would measure the number of alarms caused by objects that "couldn't" conceal a threat, although probably more useful than the absolute FAR measure adopted, is so subjective as to be effectively impossible to measure. Such a relative FAR would have a ideal level of 0%.

For absolute measures of PD and FAR, it makes no sense to define a separate PD and FAR for DCRs, imprecisely-defined dark regions that "could" contain a threat. Such regions are more reasonably accounted for by consideration of occluded threats, which may or may not be feasibly considered for a particular threat class. The FAR must still be evaluated subjectively, however, in cases where occluded threats can be considered in a threat class.

A FAR can be defined for any well-defined threat class; however, an overall FAR for all threats is also an important statistic. Ideally this would be confined to 10-15%. If the overall FAR (corresponding to a search recommendation) became significantly higher than 20%, say, then one would have to question the usefulness of the system, even if one could prove that all alarms are valid (e.g., indicate DCRs that could hide a gun, or, worse still, a detonator). Such practical constraints on the FAR, therefore, impose limitations on what threat classes the system can attempt to detect, and particularly which occluded objects can be considered.

THIS PAGE LEFT INTENTIONALLY BLANK

3.3.3 System Throughput

The desired system throughput was set at 12 bags per minute, or an average of 5 seconds per bag. A typical bag takes 2.6 seconds to be scanned, so, on average, only an additional 2.4 seconds are available after all data are available.

If data collection and image processing are overlapped, or "pipelined", then a full 5 seconds, on average, are available for processing. The use of a digital interface makes this fairly straight forward.

3.3.4 Basic System

It was decided that the basic system would concentrate on detecting handguns and dark regions that could hide handguns in real-time. The latter would also include regions that could hide other dense threats such as pipe bombs.

Four classes of handguns, as shown below, were considered.

Table 3.3 Classes of Handgun Images

- Flat guns with the handle visible
- Guns with the barrel visible but not the handle (barrel aligned horizontally)
- Guns with the barrel aligned vertically
- Guns occluded by, or touching, other dense objects

Classes 2 and 3 are generally referred to as "non-flat" guns. Separation of guns into the four classes was motivated by the realization that the algorithms would use different criteria in each case and the resulting performance would likely differ significantly. Occluded guns can never be a very well-defined class, however, and any results for this class will necessarily be heavily dependent on the sample employed.

3.3.5 Advanced System

It was decided that the aim of the advanced system would be to improve the gun detection capabilities of the basic system and to add the capability of detecting

non-occluded knives and scissors.

4. SYSTEM DESIGN

4.1 Hardware Overview

The following diagram shows the various components of CAXSS. The touch screen is for the advanced system only.

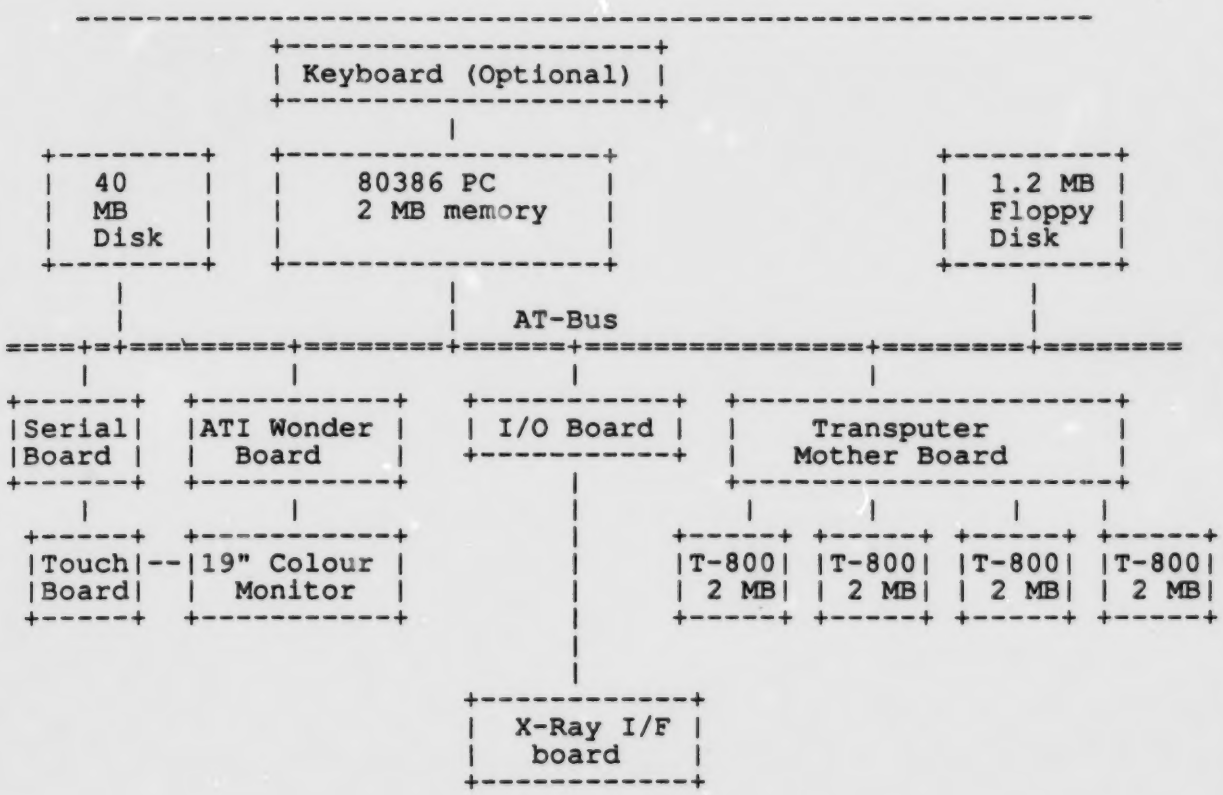

Figure 4.1 CAXSS Prototype System Configuration

All boards in the system are housed in a standard PC chassis, with the exception of the x-ray interface board which is placed in the "zoom" slot of the Linescan machine, and connected by cable to the PC.

Four transputers are shown, although only two are used in the basic system. One of the transputers, designated the "master" transputer, is responsible for communications with the 80386 via the AT bus. A communications process on the master transputer routes messages and data to and from the other transputers.

The ATI Super-VGA Wonder Video Card was considered to be the best display board on a price-performance basis, although two other display boards (the MATROX MVP-AT and the Everex Vision-8) were also tried and are supported by the system. The ATI board provides 480 x 640 x 8-bit resolution. A maximum size x-ray image consists of 480 x 512 x 8-bits, which leaves sufficient space for a touch-screen panel on the screen.

No board-specific capabilities are used so any other memory-mapped display board could be easily employed, although boards that do not support both text and graphics modes must be used with a separate text monitor.

4.1.1 Digital Interface to Linescan One

The digital interface has two purposes: to reduce the noise effects of additional analogue-to-digital and digital-to-analogue conversion, and to buffer data to the PC so that the first regions of an image can be processed while later regions are still being compiled. Regions comprise 64 lines of 480 pixels (a total of 30720 bytes). The time needed to transfer this amount of data from the x-ray machine to the buffer board is approximately 0.5 seconds. Once a region has accumulated in the buffer board, it is transferred to the PC memory via a DMA (Direct Memory Access) channel. This transfer takes only about 25 ms. This leaves 475 ms / region for processing.

The interface consists of two custom-designed boards: one (the "line driver board") resides in the zoom slot of the Linescan unit; the other (the "buffer board") resides in the PC. Connection between the two boards is via a cable consisting of 8 differential data bits, 4 differential control lines and a cable to power the circuits on the line driver board.

Astrophysics now offers a newer model of x-ray machine, Linescan System Two. The "input board", which converts the analogue x-ray signal in the diodes to digital form, has been updated. The earlier version uses an analogue compensation technique, while the newer one uses digital methods at 12-bit resolution. With the newer board, compensation for individual diode variances is greatly

improved, thus giving a superior enhanced image with less density variation from diode to diode, and more consistency over time.

The custom digital interface was designed for the Linescan System One, which uses an array of 480 diodes to scan a line across its conveyer belt; Linescan System Two uses 512. The method of "grabbing" digital data from the x-ray system bus is not compatible with 512 pixels per line. The possibility of making the interface board capable of automatically detecting which input board (One or Two) is being used and operating correctly with either was considered, but rejected because a major redesign of the interface would be required. Compatibility can be achieved much more easily with manual input board selection (setting two jumpers on the buffer board and altering a pin connection on System Two input boards); this results in 480 pixels per line coming from the System Two board.

4.2 Software Overview

4.2.1 Introduction

CAXSS software runs on the PC and on the transputers. In general, the PC is responsible for data acquisition and display, the operator interface, and overall process control. Most image analysis is done on the transputers. The raw data is displayed as soon as it is received; threats detected by the transputers are highlighted on the screen as their locations are reported to the PC.

Some region-based preprocessing for guns takes place on the PC. All processing for blade detection is handled by one dedicated transputer.

4.2.2 Input from X-ray Machine

The x-ray system scans a line across its conveyer belt via an array of 480 diodes, thus producing a 480-pixel image line. A typical bag consists of about 320 lines of data and a very large bag might use up to about 500. An upper limit of 512 lines was adopted for display and processing purposes.

The data are passed to the PC subsystem from the x-ray interface in blocks of a fixed number of lines (64). As a block is received, a DMA operation is initiated to transfer the raw data to the master transputer, while the PC simultaneously displays the data on the colour monitor. The master transputer passes the raw data on to the other transputers.

4.2.3 Processing on PC

The PC preprocesses the incoming blocks (regions) of raw data to find DCRs (Dark Connected Regions) that could possibly be guns. This is done in parallel with acquisition of the rest of the image. To find a DCR, the PC uses a simple region-growing algorithm based on an absolute threshold. In some cases, a DCR in the current data block has to be merged with DCRs in previous blocks. As soon as a complete DCR has been segmented, the PC frames it on the display.

Then the PC makes a preliminary assessment of the threat potential of the DCR. If the DCR is sufficiently large and dense, it is assigned a yellow threat level and outlined in yellow on the display. No further processing is done on yellow threats.

If a yellow threat level is not assigned, a DCR Request message is sent to the transputers, identifying the position of the DCR in the raw data. No DCR Request message is sent for yellow threats; the purpose of this is to avoid the great amount of processing time required to analyze a very large object.

After image data collection is complete, the PC waits for the transputers to finish their assessment of any remaining DCRs and other threats, and highlights the image accordingly as Threat Descriptors are received. During this time, processing may be aborted by the reception of new data from the x-ray system. If this occurs, the PC informs the transputers, which abort and prepare to process the next image. This approach was chosen on the basis of simplicity and was considered adequate for a prototype. In a real operational system, data collection and processing clearly should be "pipelined", so that data from a following bag can be collected while an earlier bag is still being processed.

4.2.4 Processing on Transputers

The main transputer is responsible for distributing data and orders to the others. It is also able to do threat analysis itself. As previously mentioned, one transputer is dedicated to blade analysis. Each of the other transputers contains a complete gun identification process so that they can work completely independently of each other. The master transputer sends the current DCR request to the next available transputer when there is one free to process it. If processing by a transputer results in a gun outline that overflows the original DCR into another one, any redundant pending requests are cancelled.

The distribution of DCRs in real airport imagery is such that four transputers are effectively utilized with this simple scheme. Since the use of more than four transputers was not considered cost-effective, there was no apparent benefit to parallelizing the processing of a DCR over more than one transputer. This would also have the disadvantages of significantly increasing the complexity and decreasing the portability of the software.

The threat analysis software on the transputers is described in more detail in Chapter 5.

4.3 User Interface

The users of the system are classified into 3 distinct levels as shown below.

Table 4.1 CAXSS User Classifications

. Operator :	Regular every-day users of the system.
. Supervisor :	A more sophisticated user, usually in charge of the operator(s).
. Analyst :	A system analyst, who may perform system installation, maintenance or debugging.

Each level has access to all the privileges of the level(s) above, and more. In general, references to "user" will include operators and supervisors only.

The user interface is based on the concept of a window, shown as a rectangular outline on the screen. The window may be moved to any position, and its size modified as required. A window is automatically positioned by CAXSS to surround the most dangerous object (as determined by the image processing algorithms). If no threat is detected in the image, a large window is placed around the entire bag, in case the operator wants to examine the bag more carefully. The user can perform various operations within the window, such as zooming, histogram equalization and other kinds of filtering. He/she can also enable a false-colour display, which enhances the density information in the image. The user can select a "next" potential threat by moving the window to some other object. No automatic processing is done on user-defined windows.

An extensive on-line help facility is available, with a choice of either a brief single help screen, or a detailed set of screens describing each function. For analyst use, there is a system status display on the PC monitor.

The system, by default, runs in operator mode. In order to invoke supervisor mode, a user must enter a valid password. Another password is needed to move from supervisor to analyst mode.

A user may use either the keyboard or the touch screen to interact with the system. The touch-screen interface consists of several icons (an icon being a rectangle containing a picture that implies a particular function). Each icon is equated with a programmed function key; touching an icon is equivalent to pressing the associated key. All supervisor-accessible menus are also configured to allow selection of menu items by touch. This allows the system to be fully operational without the need for a keyboard to be attached; this is the intended mode of operation.

4.3.1 Operator Interface

The Operator interface allows the user to manipulate the image on the display screen via interactive commands. The key definitions are shown below.

Table 4.2 Operator/Supervisor Functions

	Key	Function
*	F1	Macro key 1
*	F2	Macro key 2
*	F3	Macro key 3
*	F4	Macro key 4
*	F5	Cycle window to next threat
*	F6	Grow window
*	F7	Zoom up within window
*	F8	Clear all highlighting
	F9	Apply histogram equalization look-up table within window
	F10	Apply Pseudo-colour look-up table
	Shift-F1	Brief Help
	Shift-F2	Detailed Help
*	Shift-F3	Median filter within window
*	Shift-F4	Toggle highlighting within window
	Shift-F5	Display histogram within window
*	Shift-F6	Shrink window
	Shift-F7	Zoom down within window
	Shift-F8	Display highlighted image
*	Shift-F9	Apply normal look-up table
*	Shift-F10	Invert look-up table within window
*	Cntl-F7	Reset Zoom
*	Cntl-F9	Cycle through: Normal/Histogram Equalization/ Pseudo-colour
*	Alt-F1	Enter supervisor mode (password required)
*	e	Archive image
	Arrow Keys	Move window in direction indicated

The keys marked with an asterisk are mapped to appropriate icons on the touch screen. The touch screen allows moving and resizing a window in a natural manner. The user can loop through the threats identified by the system by cycling the current threat window, or define a new window with which to examine some other object.

A customized selection of these functions is enabled by the supervisor for use in operator mode. Only the icons for enabled functions are display on the touch screen.

The help functions (Shift-F1 and Shift-F2) switch the screen to text mode temporarily. Upon return, the graphics screen is restored.

Function keys F1 - F4 are the supervisor-defined macro keys. The functionality of a function key is definable by the supervisor to be any sequence of valid supervisor functions, or none at all, in which case the key is effectively disabled.

The default functions are shown below.

Table 4.3 Default Operator Macro Functions

Key	Function
F1	Move window to next object and highlight
F2	Display highlighted image
F3	Zoom and perform histogram equalization within window
F4	Zoom and median filter within window and pseudo-colour image

---

Preliminary investigation has found the functions defined above for F3 and F4 to be very useful in the interpretation of dark regions. Median filtering results in more uniform pseudo-colouring that is relatively easy to interpret.

4.3.2 Supervisor Interface

The Supervisor Interface allows access to all of the basic functions. The supervisor can also map any combination of commands to the four definable function keys which the operator may use.

The additional keys available in supervisor mode are shown below.

Table 4.4 Additional Supervisor Mode Functions

Key	Function
Alt-F1	Enter Analyst mode (password required)
Alt-F2	Enter Operator mode
Alt-F3	Supervisor control menu
Alt-F4	Display detailed statistics for object in window

In addition, all keys defined for operators are available in supervisor mode, regardless of whether they are disabled in operator mode.

Alt-F4 results in the display of object statistics on an optional text monitor, generally to be used by analysts only. The sort of statistics presented depends on what kind of threat is currently in the threat window.

Alt-F3 will cause entry to a supervisor menu of control functions (as opposed to processing functions which are mapped to keys as shown above). The primary options are shown below.

Table 4.5 Supervisor Control Menu Options

- Change operator macro key definitions
- Define initial key sequence
- Enable/Disable operator keys
- Calibrate system
- Change supervisor password
- Change audible alarm

In order to define the operator macro keys (F1-F4), the supervisor must enter a sequence of valid supervisor keys. The keys in the sequence do not need to be individually enabled for use by the operator. All key changes are permanently recorded in a disk file.

The supervisor can also specify an initial key sequence that will be automatically applied to the highest level threat identified in the image. Typically this might be something like histogram equalization, possibly followed by zooming up.

The supervisor has the ability to enable or disable any of the system functions for operator use. If desired, all functions may be disabled, or only one or two of the macro keys might be enabled. Only the icons for enabled functions are displayed on the screen when the system is returned to operator mode. Giving supervisors the ability to restrict operator functionality in this way is considered important by operations personnel.

The supervisor can configure the alarm to sound never, only for red level threats, for both red and yellow, or for all perceived threats - red, yellow and blue.

4.3.3 Analyst Interface

The Analyst interface is intended only for system developers. Typically, a text monitor is configured into the system for use in analyst mode. Diagnostic information can then be displayed on the text monitor during processing and while the image is being displayed and highlighted on the colour graphics monitor.

The analyst has access to all features of the system including: debugging facilities, system parameters, performance measuring and development aids.

a) Debugging Facilities

Extensive debugging facilities are built into CAXSS which allow an analyst to track down system problems quickly and accurately.

The debugging screen is displayed on the text monitor. It consists of several windows. There is a window for each transputer, where debugging messages are displayed. The two most important interfaces, from the X-Ray machine to the PC, and from the transputer to the PC, are monitored by several status indicators. As PC memory is dynamically allocated, the status of the buffers are displayed on the screen. As each image processing step on the PC is completed, a status indicator is set. A scrolling logo is very useful in indicating if the system has 'hung'.

b) System Parameters

The analyst is able to modify system parameters that

enable/disable analyst features or affect system performance and detection. The system parameters are divided into several categories:

- PC Image Processing Parameters
- PC Control Parameters
- Transputer Control Parameters
- Transputer Gun (Dense Threat) Parameters
- Transputer Blade (Knife) Parameters

Every parameter may be modified, using the keyboard, from within CAXSS. The various parameters within each category are described in Appendix B. A reading of this appendix gives a good idea of the many features and versatility of CAXSS.

c) Development Facilities

Various additional image analysis facilities are also available for use in algorithm development, which are not generally useful or meaningful to operations personnel. These include displaying intensity profiles along lines in the image, thresholding and applying other look-up tables.

The parameter system was set up so that any of the parameters may be modified by an analyst and the current image may be re-processed using the updated parameters. For example, this allows the effect of varying a parameter to be studied in a very efficient manner.

The image archiving feature can save on disk either the raw image or whatever is on the display. This has several uses. First, test images can be saved while the system is in an operational setting (such as an airport), for later experimentation and/or assessment. Input to CAXSS can be from disk files instead of from the x-ray interface. Second, images processed by outside packages and then input to the system can be used to evaluate new preprocessing methods quickly. Third, this feature can be used to try out preprocessing methods that are available only in the operator interface (for manually examining windows), without the effort of integrating them on a trial basis into the automatic processing stream.

Other specialized features, such as one to study perimeter chain-codes of objects, have proven useful.

CAXSS maintains a log file on hard disk which records the start-up and shut-down times of all system use. This log file is also used to record changes in the normalization parameters due to the auto-calibration

process.

5. THREAT RECOGNITION ALGORITHMS

5.1 Introduction

This chapter describes the procedures used by CAXSS to detect and classify threats. Handguns and other dense threats are the primary class addressed in the basic system. Blades and Electronics are added in the advanced system.

Image processing techniques that are used in the system and referred to in this chapter are:

- Application of Discriminant and Heuristic Functions
- Thresholding
- Convolution with Gradient Operators
- Dilation
- Erosion
- Eight-connected Region Growing
- Perimeter Chain Coding (Reference 1)

A brief description of each is given in Appendix C.

5.2 Handguns

Processing for guns is performed in three stages:

- Object Segmentation
- Feature Extraction
- Object Classification

5.2.1 Object Segmentation

One of the most computationally intensive activities is segmenting the image into distinct objects. A large amount of effort was devoted to developing a good segmentation technique. It was initially anticipated that some sort of adaptive thresholding technique would be employed. However, experimentation with adaptive thresholding techniques resulted in poor segmentation and excessive processing time.

Initial gun segmentation is done very roughly on the PC, using a simple region-growing algorithm based on an absolute threshold. It works on bit-packed thresholded data employing a complex technique that primarily uses boolean operations on bit-packed words. This pass is intended to identify all possible objects of interest in

an image. Some care is also taken to avoid merging distinct objects. However, it is possible for the PC to find separate dense parts of the same complex objects twice or more with this naive algorithm. Overlap is dealt with by the transputers. The DCRs segmented out in this pass are examined more closely by the transputers.

The best technique identified, which is used in the transputers, is region growing using inclusion tests that involve a gradient operation and two thresholds, referred to as the core threshold and upper threshold. The core threshold is a relatively low value that corresponds to the pixel levels of only the densest parts of a gun. The upper threshold is a level below which most parts of a gun fall. These values are determined by use of a calibration shim initially, and then adaptively varied to maintain a constant false alarm rate (described later).

If eight-connected region growing is performed using the upper threshold, then, typically, most of the mass of a handgun is found. However, the handle of a revolver, which connects to the frame via thin strands of metal, is frequently missed, as is the trigger of most guns. The edges of the gun are also poorly defined. Typically, the edges of the gun including the trigger and handle connections of a revolver are high-gradient points. The problem is solved by adding a second criterion for inclusion in the region: an otherwise-rejected point is accepted if if it is adjacent to a previously accepted gun point and if it has a sufficiently high gradient value.

Originally a Sobel operator was used for the gradient, but it was slow. The final version uses a simple four-point Robert's cross. The degradation in performance is negligible and processing is much faster. When growing is performed using a suitable gradient threshold, almost all guns can be fully segmented.

This segmentation approach has a definite weakness in that the use of the high-gradient points make it sensitive to occlusion by thin objects. For example, when a gun touches the edge of a dense briefcase, the system may include part of the briefcase in the gun outline, in some instances. Fortunately, this does not seem to interfere significantly with detection. Attempts at correction have not been totally successful; the most effective methods increase processing time far beyond an acceptable level. A series of erosions followed by dilations works reasonably well in many cases.

a) Before segmentation

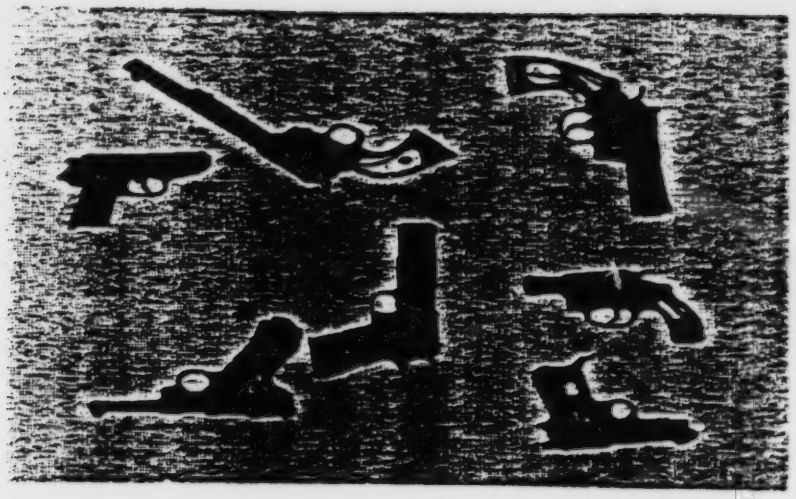

b) After segmentation

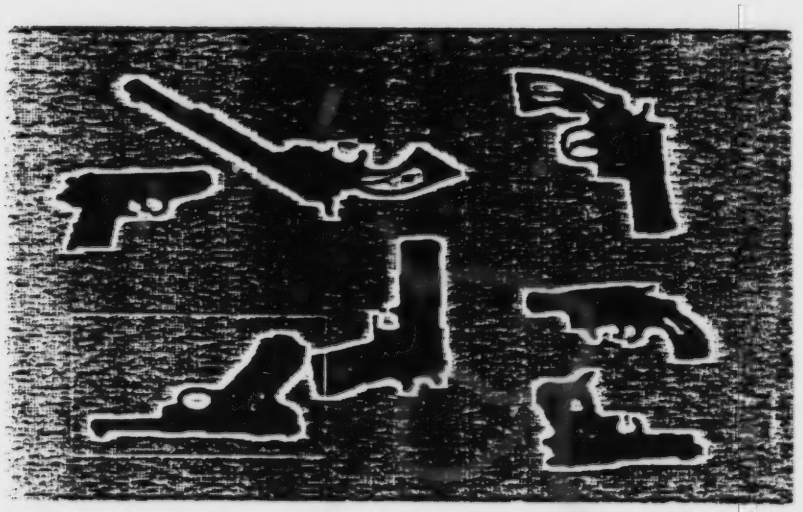

c) Barrels and Handles

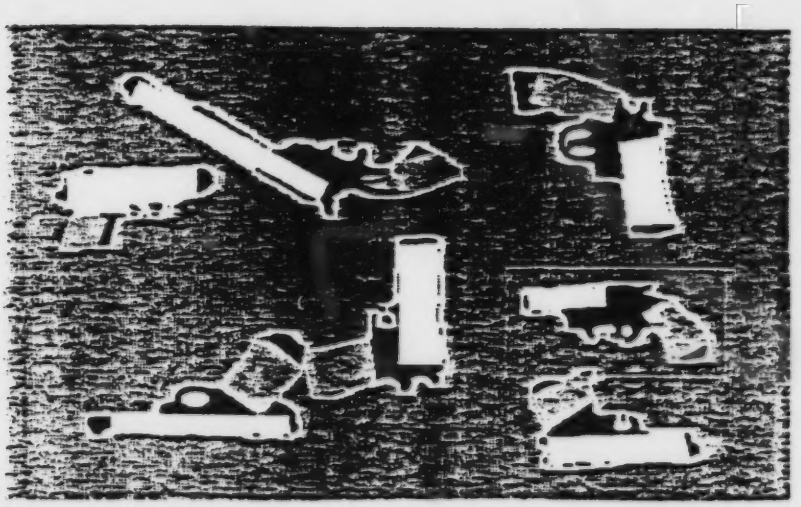

Figure 5.1 Flat Non-occluded Guns

The segmented object is now thresholded with the core threshold. The resulting "core points" are dilated. Then region growing is performed to identify the largest set of connected core points, termed the "largest core". This typically represents the solid frame of a handgun.

Finally, holes in the object are filled in. The background level around the object is sampled and used to determine a threshold below which any internal points are considered to be part of the object. This is done efficiently by bit-packing the object and employing an approximate algorithm that considers a point to be internal if there also exist object points above, below, to the right, and to the left of the point.

The results can be seen in Figures 5.1 and 5.2.

Many variants of this segmenting procedure were attempted and rejected. Initial filtering of the image was found to provide some benefit, but not enough to justify the substantial increase in processing time. Attempts were made to impose extra homogeneity requirements on region growing. For points between the core and upper thresholds, both a low gradient requirement and a maximum difference from the local object mean were tried. At best, they were found not to degrade the segmentation, but no benefit was observed.

5.2.2 Feature Extraction

A large number of features (statistics) are computed for each segmented object. Basic features are shown below.

Table 5.1 Basic Object Features

Area	= Number of pixels
Density	= Average pixel value
Mass	= Sum of (255 - pixel value)
Variance	= Variance of pixel values
Diameter	= Number of erosions required to eliminate the object
Perimeter	= Number of points in smoothed perimeter

The density, mass, and variance are computed before hole-filling and the others are computed after hole-filling. A dilation and an erosion are performed before

a) Before segmentation

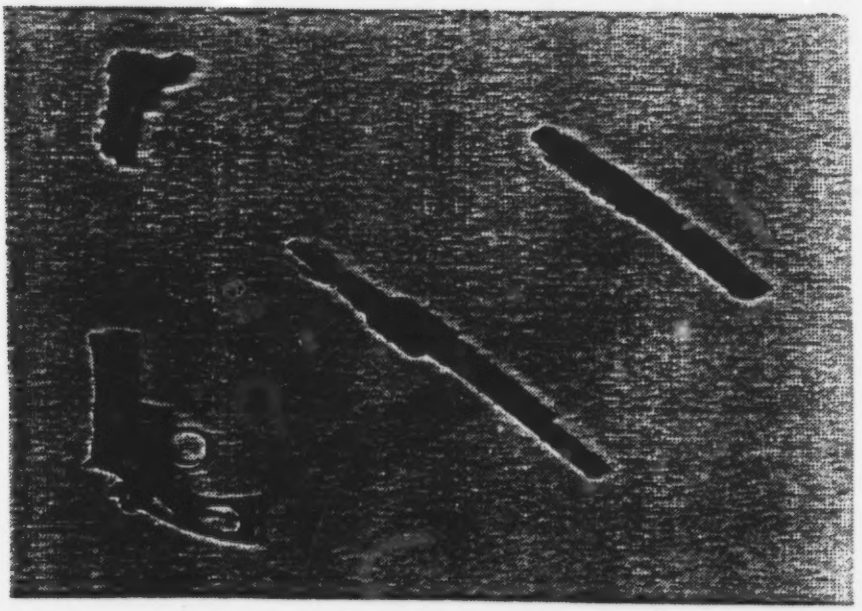

b) After segmentation

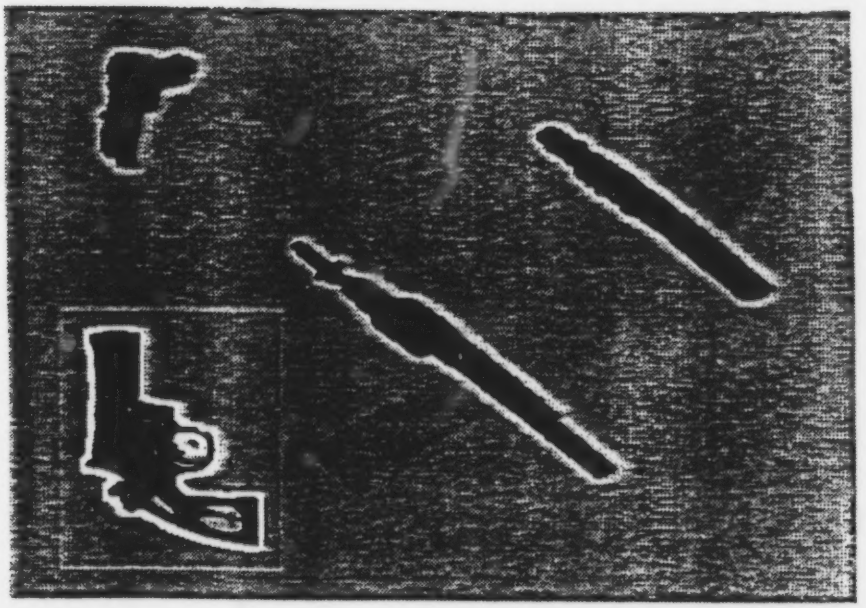

Figure 5.2 Flat and Horizontal Guns

computing the perimeter in order to smooth it somewhat. All length measurements are also corrected for the aspect ratio of the pixels, which is 4:3.

These basic features are computed both for the entire object and for the "largest core" of the object consisting of the densest pixels. The core values are particularly useful for occluded objects, since they are relatively insensitive to occlusion, except by extremely dense objects.

Another useful feature is the total number of core pixels, irrespective of whether they are connected. Various derived ratios are also employed such as diameter/area and compactness. Compactness is defined so as to have value one for a circle;

compactness = 4 * pi * area / (perimeter squared).

The remaining features require some matching of the object with definite parts of a typical gun. For recognition of gun shapes, a perimeter chain-coding technique is employed (described in Appendix C). It is used, after smoothing, to identify the longest near-straight line in the perimeter, which is hypothesized to be one side of a possible barrel. The algorithm tolerates enough deviation from a perfect straight line that it can find the edge of almost all gun barrels in typical bags. From this edge, a maximal rectangular barrel is determined. With the most probable barrel located, a potential handle is found in a similar manner, using the next longest near-straight line in the perimeter.

Examples of barrels and handles found by the system can be seen in Figure 5.1 c).

Once a potential barrel-handle combination has been found, we then compute various derived statistics such as the barrel length, width, mass, density, percent coverage of the object by the barrel and handle, and percent overlap of the barrel and handle. Experimentation with large samples of non-gun objects shows that few non-gun objects exhibit a barrel-handle with properties similar to those of typical barrel-handles found for real guns.

Many other features were investigated and rejected. The time constraint greatly restricts what techniques are feasible and also restricts their number.

The first attempt at shape recognition employed a medial-axis transform technique to derive the skeleton of the object. Various properties of that skeleton,

such as whether or not it exhibited the "L" shape characteristic of a handgun, were examined. The results are naturally highly dependent on the accuracy of the skeleton. Although this method allowed identification of guns quite well, it had an unacceptabe high false alarm rate.

Another promising algorithm was developed to try to measure how strongly the object exhibits an "L" shape. A line is fitted through the object using calculations of the principal axes and moments of the points making up the object. A perpendicular is drawn through the object's geometrical centre, splitting the object into two parts. New lines are then fitted to each of these parts. The goodness of fit of the two lines and degree to which they form an "L" are then assessed. This technique, being region-based (uses all points in the object), is somewhat complementary to the perimeter chain coding technique (uses only the perimeter of the object), and probably less sensitive to errors introduced in the region growing steps. However, chain coding produces good results by itself, and processes fewer points. Since there is insufficient processing time available for both, the "L" shape technique was not pursued.

5.2.3 Object Classification

As indicated in Table 3.2, objects are classified into three threat levels of increasing severity designated blue, yellow and red. The goal for selecting threat levels for possible guns is as follows:

Table 5.2 Threat Levels for Guns

Colour	Meaning
RED	non-occluded flat guns with clearly visible handles,
YELLOW	all other non-occluded guns, and all occluded flat guns,
BLUE	all other objects that could potentially hide a gun.

All occluded guns, no matter how highly occluded, should be identified as at least a blue level threat.

Our classification technique employs a combination of discriminant functions and heuristic functions of the extracted features. The variables used for the discriminant functions were determined via a stepwise discriminant analysis which eliminated all non-significant variables. There are three levels of tests, that result in either elimination of the object from suspicion or the assignment of a threat category or colour.

All suspicious objects are subjected first to level one tests whose purpose is to eliminate obvious false alarms; objects passing these tests are tentatively coloured blue. Barrel-handle statistics are then calculated for these blue level threats and level two tests using these statistics are applied. This results in the elimination of some objects that had tentatively been coloured blue, and a separation of the remaining objects into flat guns, non-flat non-occluded guns, and occluded guns. The non-occluded guns are tentatively coloured yellow. Level three tests, run on all blue and yellow threats, detect red threats.

A blue level classification is based on a linear discriminant function and other heuristic criteria. The first level discriminant function employs the variables shown below.

Table 5.3 Blue-Level Discriminant Features

- area
- diameter
- mass
- density
- variance
- compactness
- total core area
- largest core area
- largest core diameter
- largest core density
- largest core variance
- ratio of largest core diameter to largest core area

Objects of a sufficient size and core are also classified blue even if they are not caught by the discriminant function. These functions are defined so as to

guarantee classifying virtually all guns at least as blue, without much concern for the associated false alarm rate.

An object is coloured yellow only if it was coloured blue and it passes one of the "second-level" discriminant functions and, in addition, some heuristics. There are two second-level discriminant functions. These functions employ the barrel-handle features, which are only computed for objects that are at least blue. The first second-level discriminant function is a three-class function, which is used to discriminate non-gun objects vs. flat guns vs. non-flat non-occluded guns. Features used in this function are shown below.

Table 5.4 Yellow-Level Discriminant Features

- area
- mass
- perimeter
- diameter
- density
- variance
- compactness
- total core area
- largest core area
- largest core mass
- largest core diameter
- largest core density
- largest core variance
- ratio of diameter to area
- ratio of largest core diameter to largest core area
- barrel area
- ratio of barrel area to total area
- barrel length
- barrel width
- handle area
- ratio of barrel length to barrel width
- ratio of handle length to handle width
- ratio of barrel handle intersection to barrel plus handle area
- ratio of barrel plus handle area to total area

The other second level discriminant function is used to discriminate non-gun objects vs. occluded guns. This function employs the features shown below.

Table 5.5 Occluded-Gun Discriminant Features

- area
- mass
- perimeter
- diameter
- density
- variance
- compactness
- total core area
- largest core
- largest core mass
- largest core diameter
- largest core density
- largest core variance
- ratio of diameter to area
- ratio of largest core diameter to largest core area
- barrel area
- barrel variance
- barrel width
- ratio of barrel plus handle area to total area

The functions at this level are computed so as to produce an acceptable false alarm rate at the yellow and red levels. To be coloured yellow, the objects not only need to pass one of the discriminant functions, but also need to have a reasonable size of barrel and handle.

Finally, an object is coloured red on the basis of two heuristic functions that mainly employ the barrel-handle features and various derived features. The features used in the two heuristic functions are shown below

---

Table 5.6 Red-Level Heuristic Features

- area
- compactness
- diameter
- barrel area
- ratio of barrel length to barrel width
- handle area
- ratio of handle length to handle width
- barrel-handle intersection
- ratio of barrel plus handle area to total area
- ratio of barrel-handle intersection to barrel plus handle area
- ratio of barrel-handle intersection to handle area

A blue level threat that resembles a gun very closely is coloured red. A yellow level threat that resembles a gun, possibly with somewhat less certainty, is also coloured red. The function parameters are tuned to provide a very low red false alarm rate.

5.3 Knives and Scissors

The system attempts to identify knives and scissors by looking for a sharp blade. In the current system, one transputer is dedicated to finding blade threats.

Processing for knives is performed in four stages:

- Image Pre-Processing
- Object Segmentation
- Feature Extraction
- Object Classification

5.3.1 Image Pre-Processing

For blades, it is necessary to do some pre-processing to reduce the amount of data. This is done for each 64-line region during the 533 ms. delay before the next region arrives. First, thresholding and gradient computations are performed. A flag is set when the first non-background pixels appear, signalling the start of a new bag. For each scan line in the region, the right-most and left-most non-background locations are identified. The flag is reset at the end of the bag. Blade

detection is then done only on the area containing the bag. The pre-processing (if it takes less than 533 ms. per region) does not add anything to the overall elapsed time.

The pre-processing produces a binary representation of the data. Each pixel is set to 1 if it may be included in an object containing a blade, or 0 if it is not. The choice is based on thresholds for both pixel values and gradients.

5.3.2 Object Segmentation

Once the image of the entire bag have been received, objects are segmented using the pre-processed data. Eight-connected region growing is used to extract each object in the bag.

A minimum size requirement helps eliminate noise and small objects.

5.3.3 Feature Extraction

Once a suitable object has been segmented, the perimeter of the object, after dilation, is converted into a chain-code representation, described in more detail in Appendix C. For each point along the perimeter, a directional code is calculated, with a value between 1 and 8. The code represents the direction from the previous (counter-clockwise) point along the perimeter to the current point.

The chain code is scanned for a sharp change of direction (e.g. 6 to 4 to 2) over a small number of pixels. For example, consider the perimeter of a knife pointing down. Along the perimeter, in a clockwise direction before the tip point is reached, the chain code is predominantly 6 (south). Along the tip for two to five points (depending on the segmentation), the direction is 4 (west). After the tip point, the chain code is 2 (north).

If such a change is found, then the position of the tip point can be derived. The complete tip of the blade is identified next. The chain-code is traversed from the tip point in both directions in parallel. For the typical blade, the blade tends to grow outward to form a large acute angle at the tip. The tip is considered to end where the distance between corresponding points on the two sides stops increasing. These two points and the tip point define a triangle that represents the tip of the blade.

a) Raw Image

b) Blade Analysis Results

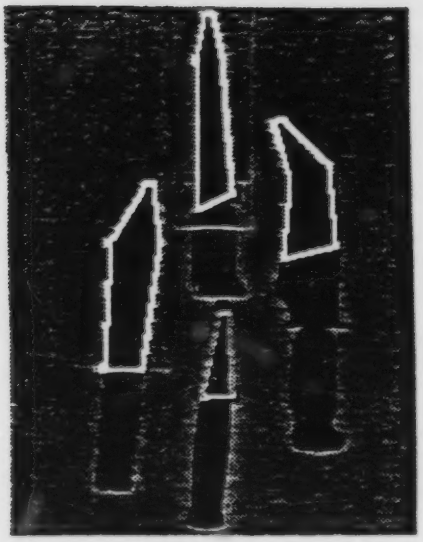

Figure 5.3 Non-occluded Knives

If a suitable tip is found, then identification of the rest of the blade proceeds. The rest of a typical blade is usually fairly rectangular. The chain-code is used to traverse the two opposite sides until the rest of the blade has been found.

After the blade has been extracted, features such as length and width are computed. The density and variance of the pixels comprising the blade are computed. These two measures alone are responsible for eliminating many false alarms in real airport bags because of the homogeneity of typical knife/scissor blades.

The area in the object around the tip is examined to ensure that it approximates a triangle. This eliminates false alarms that have rounded edges.

The features used to determine whether an object could be a blade are shown below.

Table 5.7 Blade Recognition Features

- number of points in chain-code of the tip point
- tip point width
- length of the sides of the tip (left and right)
- blade width
- length of the sides of the blade (left and right)
- density of the blade
- standard deviation of the density of the blade
- tip triangle area
- area of tip within triangle
- area outside triangle but inside the tip

The results can be seen in Figures 5.3 and 5.4.

5.3.4 Object Classification

The above features are used to classify an object into one of the three threat levels: blue, yellow or red. An object is eliminated if it does not satisfy the following requirements:

a) Raw Image

b) Blade Analysis Results

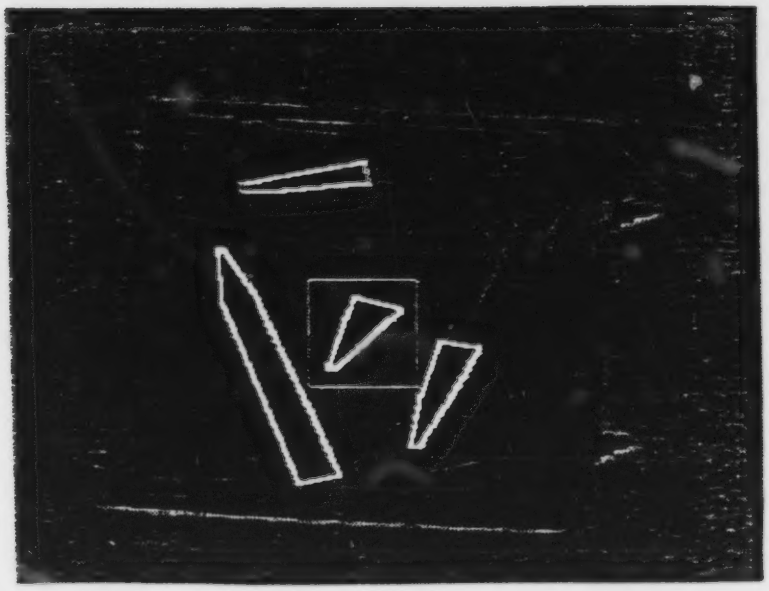

Figure 5.4 Blades in Briefcase

. minimum tip and blade length and width
. tip width not too small
. low variance compared to density in blade
. smaller variance for smaller blade lengths
. very triangular tip

The goal for discriminating between threat levels for possible blades is as follows:

Table 5.8 Threat Levels for Blades

Colour	Requirements
RED	. sufficient density . sufficient blade length and width . sufficient difference between tip and blade width
YELLOW	. same as for RED except third requirement somewhat relaxed
BLUE	. all other objects that could potentially be blades

Blue threats include blades that are short, have a low density or have a small tip-width to blade-width difference.

5.4 Other Threats

Some work has been done on the detection of electronic devices and, to a lesser extent, batteries.

5.4.1 Electronics

The current approach to finding electronic objects is based on the identifiable texture that can typically be observed in x-ray images of such items. This is characterized by substantial variations in intensity over a few pixels throughout an electronic area. For efficiency reasons, no attempt is made to segment these objects. The image is broken into rectangular cells and various measures are computed on each cell. Cells with a high-degree of electronic texture are deemed to be possible electronics. These cells, or groups of them are then circled to identify the approximate location of a

possible electronic device.

Since the x-ray image intensity of electronic boards may vary significantly from dark to light, intensity-insensitive texture analysis techniques are used. Many texture analysis techniques, such as gradient and relative extremum density methods, are local operators which effectively study the microstructure of the images and are independent of any wide dynamic range in absolute intensity.

Texture analysis of a cell is done in two stages. In the first, a gradient measure is computed, gradient features are extracted, and then a heuristic function is used to classify the cell as either non-electronics or potential electronics. The features used in the heuristic function include: average gradient value, number of relevant gradient points, absolute gradient value, and absolute intensity grey scale value to eliminate background points.

Cells classified as potential electronics by the first step advance to the second. The number of extrema in each potential electronics cell is counted. The ratio of the number of the relevant gradient points to the number of extrema in that cell is calculated. Any cell with a sufficiently high ratio is classified as electronics.

5.4.2 Batteries

A brief preliminary study was made of the feasibility of detecting batteries, using very simple techniques. The basic algorithm for identifying gun barrels was used, with strong restrictions on the absolute and relative lengths of the rectangle that was found. By measuring how well the fitted rectangle matched the object, we attempted to classify the object as a battery. This work was not pursued beyond a few test images.

6. RESULTS

6.1 Threat Detection

6.1.1 Guns

To test the system's detection of guns, a set of test images was constructed to measure the probability of detection (PD). A set of 1074 airport bags was examined to measure the false alarm rate (FAR); airport bags do not generally contain guns.

A total of 16 handguns was employed including 6 revolvers and 10 autoloaders. These guns were x-rayed in various orientations - flat, barrel-up (vertical), handle-up (horizontal), inclined 45 degrees, at the centre or edge of the belt - inside various bags and outside bags, on two x-ray machines. Sample x-rays were also taken of guns occluded with dense objects such as pieces of metal, other guns and a soldering gun, or touching other dense objects.

The measures given in the table, for each category of gun at each threat level, are:

percentage = percentage of that category of gun marked with that colour

PD = probability of detection at that threat level or higher

efficiency = log(FAR) / log(PD)

An efficiency of 22% (for PD = 90% and FAR = 10%) is the goal.

Table 6.1 Gun Detection and False Alarm Results

Category	Sample Size	Threat Level			
		None	Blue	Yellow	Red
Flat Guns	363	0	12	12	339
percentage		0%	3.3%	3.3%	93.3%
PD			100.0%	96.6%	93.3%
efficiency			>99.9	68.3	63.8
Horizontal Guns	167	1	7	132	27
percentage		0.6%	4.2%	79.0%	16.2%
PD			99.4%	95.2%	16.2%
efficiency			>99.9	48.1	2.4
Vertical Guns	80	15	26	36	2
percentage		18.8%	32.5%	45.0%	2.5%
PD			81.2%	47.5%	2.5%
efficiency			7.3	3.2	1.2
Occluded Guns	111	0	8	75	28
percentage		0%	7.2%	67.6%	25.2%
PD			100.0%	92.8%	25.2%
efficiency			>99.9	31.6	3.2
Airport Bags	1074	841	132	88	13
percentage		78.3%	12.3%	8.2%	1.2%
FAR			21.7%	9.4%	1.2%

Our performance goals were achieved at all levels for flat guns and at blue and yellow for horizontal and occluded guns. It should be noted, however, that the performance on occluded guns is highly sample-dependent.

The red level is highly tuned to detect flat guns, finding over 93% of the flat guns with only a 1% false alarm rate. Most of the occluded guns that were coloured red were those guns where the occlusion did not completely hide the shape of the gun. The 7% of flat guns that were not coloured red were mostly guns where the segmentation process failed to fully segment the handle of the gun. The barrel and handle were correctly identified in virtually every case where the segmentation was complete.

The detection of horizontal non-flat guns at the yellow

level of 95% with a 9% false alarm rate appeared to meet or exceed the capability of an ideal human in our opinion. In most cases, the gun barrel was also identified correctly.

The detection of occluded guns was also very good, and is probably comparable to an ideal human at the yellow level.

For vertically-oriented guns, we were only able to detect about half at yellow and red, and about 80% over all three levels. Most of these were quite small and most were not identifiable as guns by a human.

No statistical differences were seen in the results for different x-ray machines, nor between revolvers and autoloaders. The sample included some images that had been used to develop and tune the algorithms; however, the test results on sets of completely new images were virtually identical. No systematic differences were seen between the results on guns inside bags vs. outside bags, or on guns in bags of differing darknesses.

It was observed that some flat guns were clearly more difficult to detect than others. In particular, one small revolver that was used was consistently coloured yellow and sometimes blue when it was flat and clearly visible. The reason for this was the inability of the system to segment the handle, which was connected to the frame by exceptionally thin and non-dense connections.

Study of the red and yellow false alarms showed that virtually all of these were objects that could hide guns or other dangerous objects and should therefore be searched. Most of the blue alarms also appeared to be suspicious, although it was possible for a human to recognize some of these as common objects such as keys, clusters of coins, and cameras.

For non-occluded guns, the segmentation was good. However, the visible weakness of the algorithm in other cases did not seem to affect the statistics negatively. Although the outline drawn by the system in messy airport images was often clearly not around a well-defined object, the outlines consistently included suspicious regions. Segmentation for occluded guns was often faulty. For example, a gun touching the edge of a briefcase was frequently connected to the briefcase edge. Such objects were still generally detected at the yellow level, however.

The adaptive hole-filling algorithm seemed to produce good results. During testing, the adaptive threshold

used to fill in holes, was increased slightly to fill in more than before. Although there were some minor negative side effects, such as a gun trigger occasionally filled in, it then correctly filled in the non-dense centre of almost all gun handles for which an outline had been found.

6.1.2 Knives

A set of 49 knives and 10 scissors were used for testing. The knives were roughly divided into long knives (blade > 3") and short knives (blade < 3"). The detection results are computed on a set of test images of the knives and scissors where they were placed inside typical bags and were not occluded by any dense objects. False alarm rates were measured on 1142 airport bags. The results are shown below.

Table 6.2 Knife and Scissor Detection Results

Category	Sample Size	Threat Level			
		None	Blue	Yellow	Red
Long Knives	233	8	47	26	152
percentage		3.4%	20.2%	11.2%	65.2%
PD			96.6%	76.4%	65.2%
efficiency			56.4	10.0	7.4
Short Knives	158	20	33	39	66
percentage		12.7%	20.9%	24.7%	41.8%
PD			87.3%	66.5%	41.8%
efficiency			14.4	6.6	3.6
Scissors	33	7	11	13	2
percentage		21.2%	33.3%	39.4%	6.1%
PD			78.8%	45.5%	6.1%
efficiency			8.2	3.4	1.1
All Blades	424	35	91	78	220
percentage		8.3%	21.5%	18.4%	51.9%
PD			91.8%	70.3%	51.9%
efficiency			22.8	7.6	4.8
Airport Bags	1142	980	84	29	49
percentage		85.8%	7.4%	2.5%	4.3%
FAR			14.2%	6.8%	4.3%

Thus, at the red level, the overall detection rate is 52% with a corresponding false alarm rate of 4.3%. At the yellow and red levels, the overall detection is about 70% with about a 7% false alarm rate. The efficiency level was consistently best at the blue level, exceeding our goal of 22 on this sample for both long knives and for all blades combined.

The tips of most large knives were correctly identified. The system occasionally had difficulty in identifying the complete blade, however, even after correctly identifying the tip. This was generally due to unusual curvature in the tip and blade that exceeded the tolerance of the blade-finding algorithm.

It should be noted that the long knife category is the only one consisting entirely of clearly dangerous

threats. It is probable that some of the short knives and most of the scissors could be allowed on an airplane in hand luggage.

Although the false alarm rate of 7% at the yellow level is reasonable, investigation of the false alarms showed that many of these were clearly not knives or other sharp objects.

Although the system did not attempt to detect occluded knives, testing showed that it did detect many occluded knives when the blade was still clearly visible. It was, as expected, very sensitive to occlusions of the blade though. Any occlusion near the tip of the blade, in particular, caused the knife to be missed, even when a human operator could still identify it as a knife.

6.1.3 Other Threats

The detection rate of electronics showed promise. Several fairly complex items (such as a "walkman" type radio) were reliably detected by the algorithm. However, it was found that many non-electronic "complex" regions were classified as electronics by the system. Initially, this was considered to be a problem, but it became apparent that many such areas were quite suspicious. For example, an area in an image containing a mass of wires would be identified. Many such areas appeared either to be possible explosive devices, or at least possible hiding places for explosive devices. It was noted, though, that false alarms were regularly caused by keys and umbrellas. We expect that these could be eliminated with further effort.

Work on electronics was discontinued in favour of knife and scissor detection, so a complete real-time version was not completed and quantitative tests were not performed. Some testing was done showing a detection rate on the order of 80 to 85% for radios, but with a "false alarm" rate of 30 to 35%. As noted, many of the false alarms were suspicious areas; others were common objects.

Battery detection was successful on a few test images containing non-occluded batteries (9-volt and C cell), with a small number of false alarms on these images. Although the naive algorithm used would certainly have not worked well on a wider range of images, it did demonstrate the potential detectability of non-occluded batteries.

6.1.4 Overall Alarm Rate

The overall alarm rate was measured on a sample of 1074 airport bags. The results are shown below.

Table 6.3 Overall Alarm Rate

Number of bags	1074	
No alarm	722	67.2%
Blue alarm	192	17.9%
Yellow alarm	108	10.1%
Red alarm	53	4.9%
Yellow + Red	161	15.0%
All levels	353	32.9%

Thus the total yellow alarm rate was 15%, which corresponds to about one in seven bags. Approximately one third of all bags had some object highlighted at some level. We would expect that the number of blue alarms that appear sufficiently suspicious upon close examination to require searching the bag would be similar to the number of yellow alarms that could be eliminated by careful examination. Thus we would expect the resulting total search rate to approximate 15%. This agrees with the percentage of bags that operators are expected to open.

The 33% total alarm rate is also reasonable, since a blue alarm simply instructs the operator to inspect the highlighted area. The system rejected more than two thirds of all bags as not containing anything potentially suspicious at all.

6.2 Processing Speed

An average bag (about 320 lines of x-ray data, acquired at a rate of 8.3 ms/line in the x-ray machine) should take approximately 2.6 seconds to pass under the diode array. The goal is to process about 12 bags per minute - about 5 seconds each.

6.2.1 Airport Bags

The processing speed of the system was analyzed in detail using a sample of 1074 images of bags randomly collected on tape from three x-ray machines in use at Pearson International Airport. These images were transferred to disk, and then input, a region at a time, at the same rate as from the x-ray interface. However, input for a new image was held back until processing of the previous one was complete. In a real-life situation, either processing of the first bag would have to be aborted, the second bag would have to be ignored, or pipelining of data collection and processing would be needed.

Two times were measured: the total processing time from the start of data collection, and the processing time after all data had collected. The latter time is relevant in determining required inter-bag spacing if pipelining is not performed. The distribution of times is shown below. Only three transputers were in use for the test; performance would have been somewhat better with the advanced system's four transputers.

Table 6.4 Processing Time for 1074 Airport Bags

	Total Time (Seconds) (3 transputers)	Time After Data Collection (Sec.) (3 transputers)
Mean	4.80	2.17
Standard Deviation	1.55	1.09
Minimum	2.25	1.08
10th Percentile	2.97	1.22
20th Percentile	3.46	1.31
30th Percentile	3.90	1.47
40th Percentile	4.23	1.70
50th Percentile	4.56	1.93
60th Percentile	4.94	2.18
70th Percentile	5.33	2.46
80th Percentile	5.93	2.76
90th Percentile	6.76	3.36
Maximum	16.48	13.82

Only one bag required in excess of 10 seconds total time. The average data collection time was 2.63 seconds

with 10th and 90th percentile data collection times of 1.6 and 3.73 seconds respectively.

The same data was used to estimate the system's achievable throughput rates, assuming that scanning of the next bag does not begin until all processing of the current bag has completed. The number of bags processed in each successive minute was counted. The average throughput rate was 12.6 bags per minute with a standard deviation of 1.4 bags using three transputers, which meets the system goal of 12 bags per minute. Below, the distribution of the number of bags processed per minute is shown.

Table 6.5 Distribution of Bags Processed Per Minute

Processing Rate (bags per minute)	Number of One-minute Intervals Achieving That Rate (3 transputers)
8	1
9	0
10	3
11	18
12	22
13	18
14	17
15	5
16	1
	85

There was only a single 1-minute interval (out of 85) in which fewer than 10 bags were processed. These results would be significantly improved by pipelining.

Processing is definitely slowed down by data collection. For example, when image were input from disk at maximum speed, the throughput rate on the same sample increased from 12.6 to 14.6 bags per minute.

6.2.2 Test Bags Containing Guns

To provide an estimate of the processing speed when a bag actually contains a handgun, a sample of 98 bags each containing one handgun was tested. The processing times were not significantly different from those for

airport imagery. The average total time was 5.1 seconds with a standard deviation of 1.1 seconds. The times ranged from a low of 1.3 seconds to a high of 6.0 seconds. The average processing time after data collection was 2.5 seconds with a standard deviation of 0.9 seconds.

The overall throughput rate was 12.25 bags per minute, ranging from a minimum of 10 to a maximum of 14 bags processed each minute, in the worst case using three transputers.

6.2.3 Bags Containing Knives and Scissors

Most of the imagery collected containing knives and scissors included multiple threats in each bag so that a reasonable sample containing one threat per bag was not available. The sample used included 129 bags containing 424 knives and scissors (an average of 3.3 threats per bag).

On this sample, using three transputers, the average total time was 6.7 seconds with a standard deviation of 1.8 seconds. The total processing time after data collection averaged 4.3 seconds with a standard deviation of 1.5 seconds. The 10th and 90th percentiles for the total time were 5.2 and 7.9 seconds respectively. The average worst-case throughput was estimated to be 9.2 bags per minute with a standard deviation of 1.1 bags per minute. The number of bags processed each minute varied from a low of 7 to a high of 10.

These bags represent an extreme worst-case for the processing time since each contained multiple objects, purposely kept separate, each of which had to be analyzed. As a result, the estimated throughput rates are not very meaningful. It is our opinion that a sample of bags with one knife in each would be processed almost as fast as the airport images.

6.3 User Interface

Although no quantitative results are available, the user interface seems to have been a major success.

The touch screen allows very easy use of the system and people can be trained to use it in a matter of minutes. The window manipulation functions and image processing functions can all be used with few touches of the screen.

The coloured highlighting was effective and usually did not have to be removed in order to examine an object.

Many of the interactive processing features were extremely useful in identifying objects. Zooming was frequently used to examine small objects. Histogram equalization was found to be very useful for dark areas of images, frequently revealing identifiable structure in areas where none was noticeable before.

Pseudo-colouring was found to be useful occasionally in examining dark areas. When used in conjunction with median filtering, dense sub-components of dark objects are readily identifiable.

During testing, a definite danger of over-reliance on the system was observed. That is, operators may tend to examine only areas of images that are highlighted by the system. A suggested solution - to randomly overlay phantom guns on the image display - was implemented and found to be fairly effective. The implementation was very simple and resulted in the gun being placed partially outside the bag on occasion, but when the gun was correctly placed in a bag, it seemed to provide a good test of an operator's alertness. No facilities were incorporated for actually monitoring the operator's performance in detecting these guns, although this was suggested as a useful capability.

6.4 Developed Software

Approximately 30,000 lines of C code were developed, about half for the PC and half for the transputers. A small amount of assembly code was also used. The software is largely hardware-independent other than the I/O code. We would expect that most of the code could be retained and ported to a different multi-processor fairly easily.

The current software supports three different display boards: the ATI Super-VGA Wonder Video Card, the MATROX MVP-AT, and the Everex Vision-8. No board-specific capabilities were used so any other memory-mapped display board could be easily employed, although boards that do not support both text and graphics modes must be used with a separate text monitor.

7. CONCLUSIONS

The system detected handguns with a high degree of success. The performance in detection of handguns exceeded the PD (probability of detection) and FAR (false alarm rate) goals for all categories except vertically aligned guns. Qualitatively, few weaknesses were observed. The system seemed to meet or exceed the capability of human observers in terms of detection at the yellow and red alarm levels for all categories of guns. Almost all false alarms observed in a large sample of airport imagery appeared to be suspicious items that should be hand searched. There does not appear to be much room for improving the overall PD and FAR, although the segmentation could be improved somewhat.

The system detected long knives quite well, meeting our overall performance goals at the blue level. Overall performance goals were also met for all knives and scissors combined. Although the false alarm rate was reasonable, many false alarms observed in airport bags were clearly not knives. There appeared to be room for significant improvement in the false alarm rate.

The system met the processing speed goal of 12 bags per minute throughput rate using only three transputers. This can be improved by using more transputers, by using faster processors and/or by pipelining data collection and processing.

We would generally conclude that the prototype that has been developed could be converted fairly directly into a very useful operational system. The system would clearly enhance overall gun detection capability in Canadian airports. We also believe that the interactive processing capabilities of CAXSS could lead to further increased detection of all threats as well as to the need to search fewer bags by hand, since they can be "searched" with CAXSS.

Over-reliance on CAXSS can be limited by randomly superimposing un-highlighted phantom guns on real images, and monitoring an operator's performance.

Preliminary work on electronics detection suggests that it is probably feasible to detect electronics and other suspicious "complex" areas with reasonable PD and FAR results.

Insufficient work was done on battery detection to make

definitive conclusions. It may be feasible to detect non-occluded batteries.

8. RECOMMENDATIONS

We recommend that a complete operational system be developed for field trials. We believe that only in this manner can the true effectiveness of the system be evaluated. The primary modification needed is to ensure that CAXSS can handle real-world throughput requirements (12 bags per minute). Although CAXSS currently averages over 12 bags per minute for typical bags, with 3 transputers, the processing of a very complex bag may take noticeably longer than average and will be aborted when acquisition for the next bag begins. This situation could be handled either by pipelining bags, or by converting the system to process a scrolling display of bag imagery. The fourth transputer should also help.

At the same time, we recommend that development continue - to improve the existing algorithms, and to develop and incorporate algorithms for the detection of additional threats. In particular, we think that the work on detection of electronic items and batteries can be continued and produce good results. It should even be possible to achieve reasonable performance in the detection of some kinds of explosive devices, with methods of detecting some components of bombs, such as wires attached to batteries.

We also recommend that the system be integrated with a dual-energy x-ray machine capable of discriminating organic material from non-organic. This would provide a critical additional component in a comprehensive bomb recognition model.

REFERENCES

1. Safaee-Rad R., B. Behabib, K. C. Smith and K. M. Ty, Position, Rotation, and Scale Invariant Recognition of Two-Dimensional Objects using a Gradient-Coding Scheme, Proceedings of the IEEE Pacific Rim Conference on Communications, Computers and Signal Processing, Vancouver, B.C., 1989, pp. 306-311.

10. GLOSSARY OF TERMS

Compactness - a computed feature used to classify objects in an image; defined to have value 1 for a circle:

$$\text{compactness} = \frac{4 * \text{pi} * \text{area}}{\text{perimeter squared}}$$

Connected core area - a dense region of an object in which each core pixel touches at least one other core pixel.

Core - the densest part of an object.

Dilation - an image processing technique; see Appendix C.

Eight-connected region - a region of an image for which each member pixel touches at least one other member, in the vertical, horizontal or diagonal directions.

Erosion - an image processing technique; see Appendix C.

False threat - a phantom gun image overlaid on a real x-ray image, but not detectable by CAXSS.

Icon - a rectangle on a touch screen containing a picture that implies a particular function.

Macro - a piece of software that programs one function key to perform a sequence of functions.

Pipeline - to allow the second of two processes operating in sequence on the same data, to begin on the first part of the data while the first process is finishing up with the last part; in CAXSS, data acquisition and preprocessing can be pipelined.

Region - used specifically in CAXSS to mean a 64-line block of an x-ray image.

Region growing - an image processing technique; see Appendix C.

BIBLIOGRAPHY

Ballard, D. H., Strip Trees: A Hierarchical Representation for Curves, CACM 24, 310-321 (1981).

Ballard, D. H., Generalizing the Hough Transform to Detect Arbitrary Shapes, Pattern Recognition 13, No. 2, 111-122 (1981).

Ballard, DE. H. and C. M. Brown, Computer Vision, Prentice-Hall, Englewood Cliffs (1982).

Brady, M. Smooth Local Symmetries and Local Frame Propagation, Proc. Pattern Recognition Image Processing, 183-191 (1982).

Brice, C. R. and C. L. Fenema, Scene Analysis Using Regions, Artificial Intelligence 1, 205-226 (1970).

Burt, P. J. and E. A. Adelson, The Laplacian Pyramid as a Compact Image Code, IEEE Trans. on Communications, 532-540 (1983).

Canny, J. F., A Variational Approach to Edge Detection, Proc. of AAAI Conf., Sept 1983.

Chellappa, R. and A. A. Sawchuk, Digital Image Processing and Analysis: Volume 2 IEEE Computer Society, Los Angeles (1985).

Chow, C. K. and T. Kaneko, Automatic Boundary Detection of Left Ventrical From Cineangiograms, Computer and Biomedical research 5, 388-410 (1972).

Chow, C. K. and T. Kaneko, Boundary Detection of Radiographic Images by a Threshold Method, in Frontiers of Pattern Recognition (S. Watanabe, Ed.), 61-82, Academic Press, New York, 1972.

Crowley, J. L., A Multiresolution Representation for Shape, in Multiresolution Image Processing and Analysis, A. Rosenfeld (Ed.), Springer-Verlag, Berlin (1984).

Davis, L. S., A. Rosenfeld, and Waszka, Region Extraction by Averaging and Thresholding, IEEE Trans. System, Man, and Cybernetics 5, 383-388 (1975).

Davis, L. S., A Survey of Edge Detection Techniques, Computer Graphics Image Processing 4, 248-270

(1975).

Davis, L. S. and A. Mitchie, Edge Detection in Textures, Computer Graphics Image Processing 12, 25-39 (1980).

Duda, R. O. and P. E. Hart, Use of Hough Transformation to Detect Line and Curves in Pictures, CACM 15, 11-15 (1972).

Frei, W. and C. C. Chen, Fast Boundary Detection: A Generalization and a New Algorithm, IEEE Trans. Comput., 988-998 (1977).

Hanson, A. R. and E. M. Riseman, Segmentation of Natural Scenes, in Computer Vision Systems (A. R. Hanson and E. M. Riseman Eds.), Academic Press, New York (1978).

Haralick, R. M., Edge and Region Analysis for Digital Image Data, Computer Graphics Image Processing 12, 60-73 (1980).

Haralick, R. M. and L. G. Shapiro, Survey: Image Segmentation Techniques, Computer Vision, Graphics, and Image Processing 29, 100-132 (1985).

Horowitz, S. L. and T. Pavlidis, Picture Segmentation by a Tree Traversal Algorithm, J. Assoc. Comput. Mach. 23, 368-388 (1976).

Hueckel, M., An Operator Which Locates Edges in Digital Pictures, J. Assoc. Comput. Mach. 18, 113-125 (1971).

Hunter, G. M. and K. Steiglitz, Operations on Images Using Quadtrees, IEEE Trans. PAMI, 145-153 (1979).

Kohler, R. A Segmentation System Based on Thresholding, Computer Graphics Image Processing 15, 319-338 (1981).

Levine, M. D. and S. I. Shaheen, A Modular Computer Vision System for Picture Segmentation and Interpretation, IEEE PAMI-3 (1981).

Marr, D. and E. Hildreth, Theory of Edge Detection, Proc. Royal Society London B207, 187-217 (1980).

Nevatia, R., A Color Edge Detector and Its Use in Scene Segmentation, IEEE Trans. Systems, Man, Cybernetics 7(11), 820-826 (1977).

Nevatia, R. and K. R. Babu, Linear Feature Extraction

and Description, Computer Graphics Image Processing 13. 257-269 (1980).

Ohlander, R., K. Price, and R. Reddy, Picture Segmentation by a Recursive Region Splitting Method, Computer Graphics Image Processing 8, 313-333 (1978)

Ohta, Y., T. Kanade, and T. Sakai, Color Information for Region Segmentation, Computer Graphics Image Processing 13, 222-241 (1980).

Pavlidis, T., Graphics and Image Processing, Computer Science Press, Rockville, MD (1982).

Pavlidis, T., A Review of Algorithms for Shape Analysis, Computer Graphics Image Processing 7, 243-258 (1978).

Pietikainen, M. and A. Rosenfeld, Image Segmentation by Texture Using Pyramid Node Linking, IEEE Trans. Systems, Man, Cybernetics 11 (1981).

Prewitt, J. M. S., Object Enhancement and Extraction, in Picture Processing and Psychopictorics, B. S. Lipkin and A. Rosenfeld (Eds.), Academic Press, New York (1970).

Riseman E. and M. Arbib, Segmentation of Static Scenes, Computer Graphics Image Processing 6, 221-276 (1977).

Rosenfeld A., R. A. Hummel, and S. W. Zucker, Scene Labeling by Relaxation Operations, IEEE Trans. System, Man, and Cybernetics 6 (1976).

Samet, H., The Quadtree and Related Hierarchical Data Structures, Computing Surveys 16, 187-260 (1984).

Samet, H. and M. Tamminen, Computing Geometric Properties of Images Represented by Linear Quadtrees, IEEE PAMI, 229-240 (1985).

Schneier, M., Calculations of Geometric Properties Using Quadtrees, Computer Graphics Image Processing 16, 296-302 (1981).

Schneier, M., Two Hierarchical Linear Feature Representations: Edge Pyramids and Edge Quadtrees, Computer Graphics Image Processing 17, 211-224 (1981).

Tanimoto, S. and T. Pavlidis, A Hierarchical Structure for Picture Processing, Computer Graphics Image Processing 4, 104-119 (1975).

Waszka, J. S., A Survey of Threshold Selection Techniques, Computer Graphics Image Processing 7, 259-265 (1978).

Wechsler, S. W. and J. Skalansky, Finding the Rib Cage in Chest Radiographs, Pattern Recognition 9, 21-30 (1977).

Witkin, A., Scale Space Filtering, Proc 7th IJCAI, 1010-1021 (1983).

Yakimovsky, Y., Boundary and Object Detection in Real World Images, J. Assoc. Comput. Mach. 23, 599-619 (1976).

Zucker, S. W., Region Growing: Childhood and Adolescence, Computer Graphics Image Processing 5, 382-399 (1976).

Zucker, S. W., R. A. Hummel, and A. Rosenfeld, An Application of Relaxation Labeling to Line and Curve Enhancement, IEEE Trans. Computers 26 (1977).

Appendix A - PROJECT SCHEDULE AND DOCUMENTS

A.1 Project Schedule

Table 11.1 Project Schedule

Task	Starting Date	Ending Date
Detailed Work Plan	May 16/88	May 29/88
Technology Review	May 30/88	Sep 6/88
Threat Definition	May 30/88	Oct 24/88
System Specification	Jun 22/88	Oct 24/88
Basic System Design	Jul 25/88	Dec 18/88
Basic System Implem.	Dec 19/88	Mar 5/89
Basic Test Plan	Dec 19/88	Feb 19/89
Basic Integ. & Testing	Mar 5/89	Apr 23/89
Basic Documentation	Apr 24/89	May 7/89
Basic Demonstration	May 8/89	May 21/89
Basic Training	May 22/89	Jun 4/89
Adv. System Design	Feb 20/89	Jun 4/89
Adv. System Implem.	Jun 5/89	Aug 6/89
Adv. Test Plan	Jun 5/89	Jul 9/89
Adv. Integ. & Testing	Aug 7/89	Oct 1/89
Adv. Training	Oct 2/89	Oct 22/89
Adv. Documentation	Sep 4/89	Oct 29/89
Adv. Demonstration	Oct 23/89	Nov 12/89
Report & Presentation	Jul 10/89	Nov 12/89

A.2 Project Documentation Deliverables

Detailed Work Plan for the Development of a Computer-Assisted X-Ray Screening System for Detection of Dangerous Objects in Passenger Carry-On Baggage, Version 1.0, June 9, 1988.

Technology Review for the Development of a Computer-Assisted X-Ray Screening System for Detection of Dangerous Objects in Passenger Carry-On Baggage, Volume I: Hardware and Literature Review, Version 1.0, August 17, 1988.

Technology Review for the Development of a Computer-Assisted X-Ray Screening System for Detection of Dangerous Objects in Passenger Carry-On Baggage,

Version 1.0, Volume II: X-Ray Part, August 31, 1988.

Threat Definition Document for the Development of a Computer-Assisted X-Ray Screening System for Detection of Dangerous Objects in Passenger Carry-On Baggage, Version 1.0, October 24, 1988.

System Specification for a Computer-Assisted X-Ray Screening System for Detection of Dangerous Objects in Passenger Carry-On Baggage, Version 1.1, December 16, 1988.

Basic System Design for a Computer-Assisted X-Ray Screening System for Detection of Dangerous Objects in Passenger Carry-On Baggage, Version 1.0, December 19, 1988.

Basic System Test Plan for the Development of a Computer-Assisted X-Ray Screening System for Detection of Dangerous Objects in Passenger Carry-On Baggage, Version 1.0, February 21, 1989.

System User Documentation for the Development of a Computer-Assisted X-Ray Screening System for Detection of Dangerous Objects in Passenger Carry-On Baggage, Version 2.0, August 10, 1989.

Advanced System Design for the Development of a Computer-Assisted X-Ray Screening System for Detection of Dangerous Objects in Passenger Carry-On Baggage, Version 1.0, June 6, 1989.

Advanced System Test Plan for the Development of a Computer-Assisted X-Ray Screening System for Detection of Dangerous Objects in Passenger Carry-On Baggage, Version 1.0, August 8, 1989.

Appendix B - CAXSS SYSTEM PARAMETERS

B.1 PC Image Processing Parameters

The PC Image Processing Parameters include:

- Thresholds for segmentation and data normalization
- Thresholds for bag separation and histogram equalization
- Number of Erosions and Dilations for segmentation
- Minimum core size for DCRs
- Minimum connected core area for early termination

B.2 PC Control Parameters

The PC Control Parameters include the following switches/settings to enable/disable various features:

- Default User Mode
- Performance Flag
- Training Flag
- Running Mode
- Delay Time
- False Image Flag
- Auto Initial Sequence Flag
- Image Archiving Counters
- Transputer Time-Out Setting
- Touch-Screen Beep
- Constant FAR (Automatic Calibration)

The system is put into the default user mode (Operator, Supervisor or Analyst) when the system is started up. The performance feature, when enabled, monitors the processing time and false alarm rates (FARs) of the system for the series of images that follows. The results are stored in a disk file for later analysis. This feature was used to collect the data for the timing and FAR results given in a later section. The training feature, when enabled, saves the features extracted for all of the DCRs in a disk file. This data is used in the discriminant analysis discussed in a later section.

The running mode parameter allows selection between different modes of operation. For example, when processing images from disk it may be desirable to insert a delay between images so that the user has time to study the images. The Delay Time parameter is used to indicate the length of the delay if the appropriate mode is enabled.

The false imagery flag is used to indicate the frequency of the addition of false threats within the x-ray imagery. A false threat is an artificial gun image superimposed on real screen data and not highlighted by the automatic system; its purpose is to force operators not to rely too heavily on CAXSS. The Auto Initial Sequence is used to enable the automatic execution of a series of interactive commands for the highest threat (if any) within each image. The Image Archiving counters are used to indicate the ID associated with an image when it is saved to a disk file. The Transputer Time-Out parameter is used to determine if and when to automatically abort the transputer, if it is taking a long time to process an image. The frequency and duration of the audible beep may be varied via parameters. The Automatic calibration parameters are used by the system to adaptively normalize the raw data (via normalization parameters), in order to achieve a specified False Alarm Rate.

B.3 Transputer Control Parameters

The Transputer Control Parameters include:

- Number of Transputers
- Level 0 Enable
- Barrel/Handle Enable
- Perimeter Enable
- Chain Code Enable
- Intersect Enable
- Threat Model Enable

The number of transputers utilized may be modified via the first parameter. It can be used to add or delete transputers without re-compiling any code. The Level 0 parameter allows the system to display, with a brown outline, DCRs classified by the transputers as non-threats. The Barrel/Handle parameter allows the system to display, for dense threats, the areas within the object that were found as the 'barrel' and 'handle'. The perimeter parameter allows the system to highlight the threats in either a perimeter or shaded-in format. The intersect parameter enables/disables the complete processing of a second threat that is found to intersect with another threat already found by the system. The threat model parameter enables/disables the combining of results for the gun and blade threat analysis.

B.4 Transputer Gun Parameters

The Transputer Gun Parameters include:

- Threshold and Gradient Parameters
- Discriminant Analysis Parameters
- Early Termination Parameters
- Barrel/Handle Parameters

B.5 Transputer Blade Parameters

The Transputer Blade Parameters include:

- Knife Detection Enable
- Threshold and Gradient Parameters
- Blade Classification Heuristic Parameters
- Blade Finding Heuristic Parameters

Appendix C - IMAGE PROCESSING TECHNIQUES

C.1 Application of Discriminant and Heuristic Functions

Discriminant analysis is used for object classification based on a number of independent numerical features of the objects. For a two-way classification, two functions are calculated, one for each class. Each function consists of a sum of terms of the form:

(feature coefficient) X (feature value)

In general, a training set of data is used to generate the coefficients. To classify an object, compute or measure the feature values, and use them to compute the values of the functions. The object belongs in the class with the highest function value.

Heuristic functions in this system are constraints on feature values that experience has shown make membership in a class more likely.

C.2 Convolution with Gradient Operators

The gradient operator used in CAXSS is the Robert's Cross, as follows:

```
| 1  0  0 |        | 0  0  1 |
| 0  0  0 |        | 0  0  0 |
| 0  0 -1 |        |-1  0  0 |
```

Convolution of an operator with an image can be thought of in pictorial terms. The centre of the operator is superimposed on each image point in sequence; at each such position, every operator element is multiplied by the corresponding image intensity, and the products summed; this sum is the gradient value for the image point. In CAXSS, the two operators shown above are convolved with the image simultaneously, and the absolute values of the individual results are averaged.

C.3 Dilation

Dilation is done to smooth the edges of a found object "core" in an image, and to fill in holes. Often, it is done before or after an erosion. Let us call the result - original core and new points - the "object".

A non-object point adjacent to a core point is added to the object if it satisfies some criteria less rigorous than those used to define the core. This process is not

recursive.

C.4 Erosion

Erosion is the opposite of dilation. Any object point with an adjacent pixel not in the object is discarded. Again, the process is not recursive.

C.5 Eight-connected Region Growing

The seed for region growing is one point known to be in the region. The object is then expanded recursively to adjacent pixels in the vertical, horizontal and diagonal directions (eight in all), if these neighbours satisfy criteria for belonging to the object. Although the idea is recursive, the implementation in CAXSS is non-recursive and uses a bounded stack.

C.6 Perimeter Chain-coding

Perimeter chain-coding is one method of representing an object in a two-dimensional image. For each point along the perimeter of the object, a directional code is calculated, with a value between 1 and 8. The code represents the direction from the previous (counter-clockwise) point along the perimeter to the current point. The codes are given below:

1 : North-East
2 : North
3 : North-West
4 : West
5 : South-West
6 : South
7 : South-East
8 : East

With these codes, it is easy to study the shape of the object, independent of its orientation. For example, a code change from 3 to 6 (NW to S) over a few pixels indicates a sharp point, and a long string of more-or-less unchanging codes indicates a straight line.

C.7 Relative Extremum Density

Relative extremum density is defined as the number of local extremum points in the object divided by the total number of points in the object. An extremum point is a point whose intensity is either greater or smaller than those of its immediate neighbours. In the case of a plateau, only one representative point is counted as the extremum. Relative extremum density is a texture feature.

C.8 Thresholding

Thresholding an image means identifying pixels that have intensity values above (or below) a given value (called the "threshold").